digital

OpenVMS Guide to 64-Bit Addressing

OpenVMS

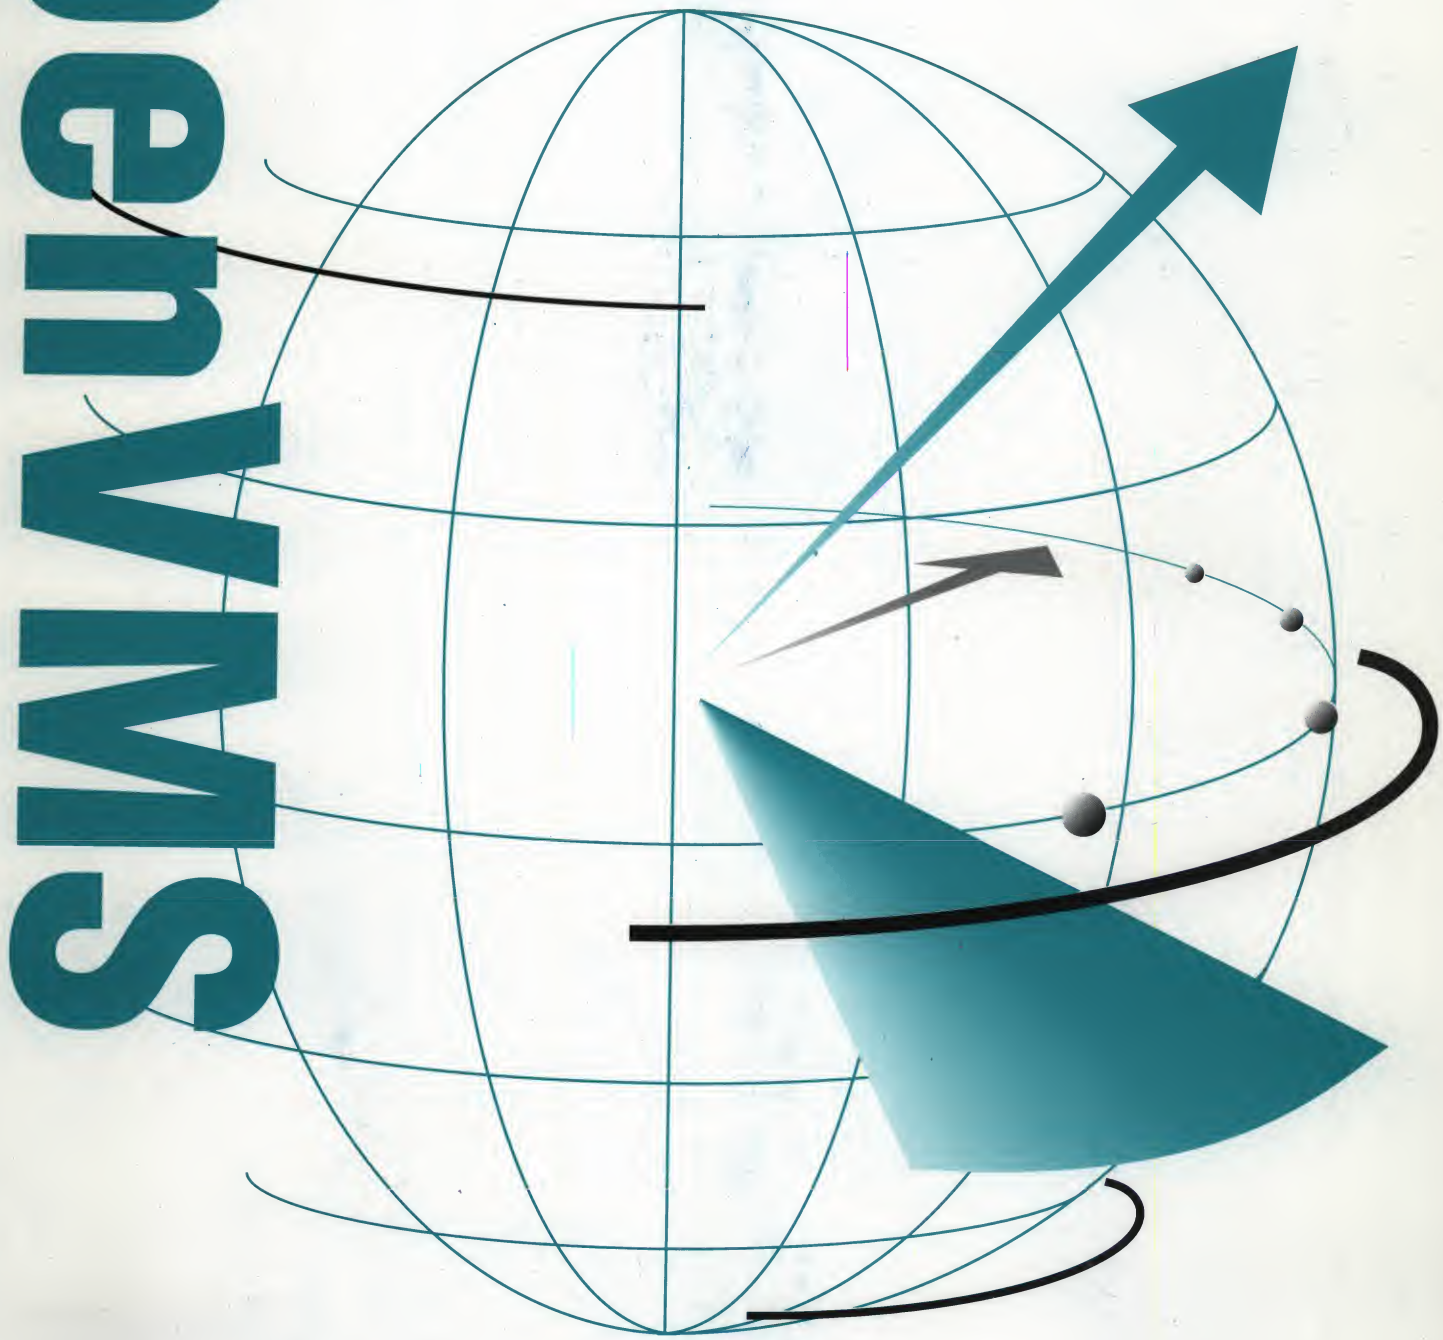

OpenVMS Alpha Guide to 64-Bit Addressing

Order Number: AA–QSBCA–TE

December 1995

This new manual describes the OpenVMS Alpha 64-bit virtual addressing support provided in OpenVMS Alpha Version 7.0.

Revision/Update Information:	This is a new manual.
Software Version:	OpenVMS Alpha Version 7.0

Digital Equipment Corporation
Maynard, Massachusetts

December 1995

ZK6467

This document is available on CD–ROM.

Contents

6 OpenVMS Alpha 64-Bit API Guidelines

7 OpenVMS Alpha Tools and Utilities That Support 64-Bit Addressing

8 DEC C RTL Support for 64-Bit Addressing

9 MACRO–32 Programming Support for 64-Bit Addressing

Preface

This guide describes OpenVMS Alpha operating system support for 64-bit virtual addressing.

The information in this document applies only to applications on OpenVMS Alpha systems; applications on OpenVMS VAX systems are not affected.

Intended Audience

This information in this guide is intended for system and application programmers. It presumes that its readers are familiar with the OpenVMS Alpha programming environment and concepts.

Document Structure

Chapter 1 presents an overview of the OpenVMS Alpha 64-bit virtual memory address space and explains how you can use 64-bit addressing in applications. The chapters that follow describe the OpenVMS Alpha programming tools and languages that support 64-bit addressing.

Related Documents

This guide provides high-level descriptions of some of the topics covered in the following manuals; refer to these books for more detailed information:

- *OpenVMS Programming Concepts Manual*
- *OpenVMS Calling Standard*
- *OpenVMS System Services Reference Manual: A–GETMSG* and *OpenVMS System Services Reference Manual: GETQUI–Z*
- *OpenVMS Record Management Services Reference Manual*
- *OpenVMS RTL Library (LIB$) Manual*
- *OpenVMS Debugger Manual*
- *OpenVMS Alpha System Dump Analyzer Utility Manual*
- *OpenVMS Alpha Guide to Upgrading Privileged-Code Applications*

  If you have an application that links against the base system image SYS$BASE_IMAGE.EXE, you might need to relink, recompile, or make source-code changes for OpenVMS Alpha Version 7.0. Refer to this manual for more information about the changes that might affect privileged-code applications and device drivers.

For additional information about OpenVMS products and services, access the Digital OpenVMS World Wide Web site. Use the following URL:

`http://www.openvms.digital.com`

Reader's Comments

Digital welcomes your comments on this manual.

Print or edit the online form SYS$HELP:OPENVMSDOC_COMMENTS.TXT and send us your comments by:

Internet	**openvmsdoc@zko.mts.dec.com**
Fax	603 881-0120, Attention: OpenVMS Documentation, ZK03-4/U08
Mail	OpenVMS Documentation Group, ZKO3-4/U08 110 Spit Brook Rd. Nashua, NH 03062-2698

How To Order Additional Documentation

Use the following table to order additional documentation or information. If you need help deciding which documentation best meets your needs, call 800-DIGITAL (800-344-4825).

Telephone and Direct Mail Orders

Location	Call	Fax	Write
U.S.A.	DECdirect 800–DIGITAL 800–344–4825	Fax: 800–234–2298	Digital Equipment Corporation P.O. Box CS2008 Nashua, NH 03061
Puerto Rico	809–781–0505	Fax: 809–749–8300	Digital Equipment Caribbean, Inc. 3 Digital Plaza, 1st Street, Suite 200 P.O. Box 11038 Metro Office Park San Juan, Puerto Rico 00910–2138
Canada	800–267–6215	Fax: 613–592–1946	Digital Equipment of Canada, Ltd. Box 13000 100 Herzberg Road Kanata, Ontario, Canada K2K 2A6 Attn: DECdirect Sales
International	—	—	Local Digital subsidiary or approved distributor
Internal Orders	DTN: 264–4446 603–884–4446	Fax: 603–884–3960	U.S. Software Supply Business Digital Equipment Corporation 10 Cotton Road Nashua, NH 03063–1260

ZK–7654A–GE

Conventions

The name of the OpenVMS AXP operating system has been changed to OpenVMS Alpha. Any references to OpenVMS AXP or AXP in this document are synonymous with OpenVMS Alpha or Alpha.

The following conventions are used in this manual:

. . .	A horizontal ellipsis in examples indicates one of the following possibilities: • Additional optional arguments in a statement have been omitted. • The preceding item or items can be repeated one or more times. • Additional parameters, values, or other information can be entered.
. . .	A vertical ellipsis indicates the omission of items from a code example or command format; the items are omitted because they are not important to the topic being discussed.
( )	In format descriptions, parentheses indicate that, if you choose more than one option, you must enclose the choices in parentheses.
[ ]	In format descriptions, brackets indicate optional elements. You can choose one, none, or all of the options. (Brackets are not optional, however, in the syntax of a directory name in an OpenVMS file specification, or in the syntax of a substring specification in an assignment statement.)
{ }	In format descriptions, braces surround a required choice of options; you must choose one of the options listed.
boldface text	Boldface text represents the introduction of a new term or the name of an argument, an attribute, or a reason. Boldface text is also used to show user input in online versions of the manual.
italic text	Italic text emphasizes important information, indicates variables, and indicates complete titles of manuals. Italic text also represents information that can vary in system messages (for example, Internal error *number*), command lines (for example, /PRODUCER=*name*), and command parameters in text.
UPPERCASE TEXT	Uppercase text indicates a command, the name of a routine, the name of a file, the name of a file protection code, or the abbreviation for a system privilege.
-	A hyphen in code examples indicates that additional arguments to the request are provided on the line that follows.
numbers	All numbers in text are assumed to be decimal, unless otherwise noted. Nondecimal radixes—binary, octal, or hexadecimal—are explicitly indicated.

1 Introduction

The OpenVMS Alpha operating system provides support for 64-bit virtual memory addressing, which makes the 64-bit virtual address space defined by the Alpha architecture available to the OpenVMS Alpha operating system and to application programs.

Many OpenVMS Alpha tools and languages (including the Debugger, run-time library routines, and DEC C) support 64-bit virtual addressing. Input and output operations can be performed directly to and from the 64-bit addressable space by means of RMS services, the $QIO system service, and most of the device drivers supplied with OpenVMS Alpha systems.

Underlying this are new system services, which allow an application to allocate and manage the 64-bit virtual address space that is available for process-private use.

By using the OpenVMS Alpha tools and languages that support 64-bit addressing, programmers can create images that map and access data beyond the limits of 32-bit virtual addresses. The 64-bit virtual address space design ensures upward compatibility of programs that execute under versions of OpenVMS Alpha prior to Version 7.0, while providing a flexible framework within which 64-bit addresses can be used in many different ways to solve new problems.

This chapter describes the layout and components of the OpenVMS Alpha 64-bit virtual memory address space and highlights the benefits of using 64-bit addresses in applications. For more information about the OpenVMS Alpha programming tools and languages that support 64-bit addressing and recommendations for enhancing applications to support 64-bit addressing, refer to the subsequent chapters in this guide.

1.1 Traditional OpenVMS 32-Bit Virtual Address Space Layout

In previous versions of the OpenVMS Alpha operating system, the virtual address space layout was largely based upon the 32-bit virtual address space defined by the VAX architecture. Figure 1–1 illustrates the OpenVMS Alpha implementation of the OpenVMS VAX layout.

Figure 1–1 32-Bit Virtual Address Space Layout

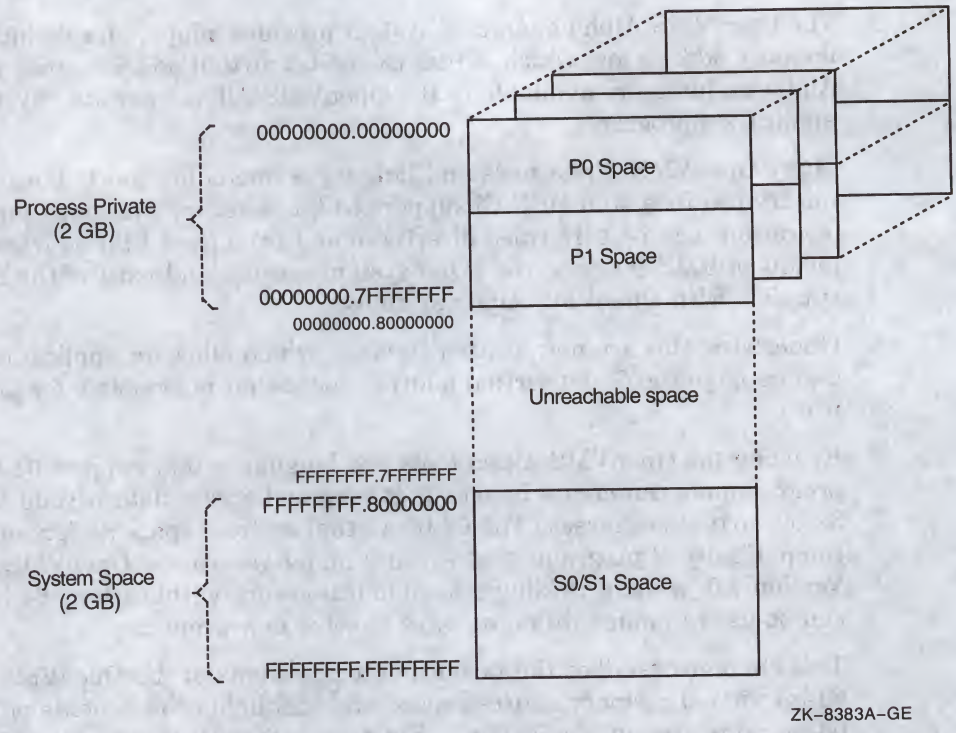

The low half of the OpenVMS VAX virtual address space (addresses between 0 and $7FFFFFFF_{16}$) is called **process-private space**. This space is further divided into two equal pieces called P0 space and P1 space. Each space is 1 GB long. The P0 space range is from 0 to $3FFFFFFF_{16}$. P0 space starts at location 0 and expands toward increasing addresses. The P1 space range is from 40000000_{16} to $7FFFFFFF_{16}$. P1 space starts at location $7FFFFFFF_{16}$ and expands toward decreasing addresses.

The upper half of the VAX virtual address space is called system space. The lower half of system space (the addresses between 80000000_{16} and $BFFFFFFF_{16}$) is called S0 space. S0 space begins at 80000000_{16} and expands toward increasing addresses.

The VAX architecture associates a page table with each region of virtual address space. The processor translates system space addresses using the system page table. Each process has its own P0 and P1 page tables. A VAX page table does not map the full virtual address space possible; instead, it maps only the part of its region that has been created.

1.2 OpenVMS Alpha 64-Bit Virtual Address Space Layout

The OpenVMS Alpha 64-bit address space layout is an extension of the traditional OpenVMS 32-bit address space layout.

Figure 1–2 illustrates the 64-bit virtual address space layout design.

Figure 1–2 64-Bit Virtual Address Space Layout

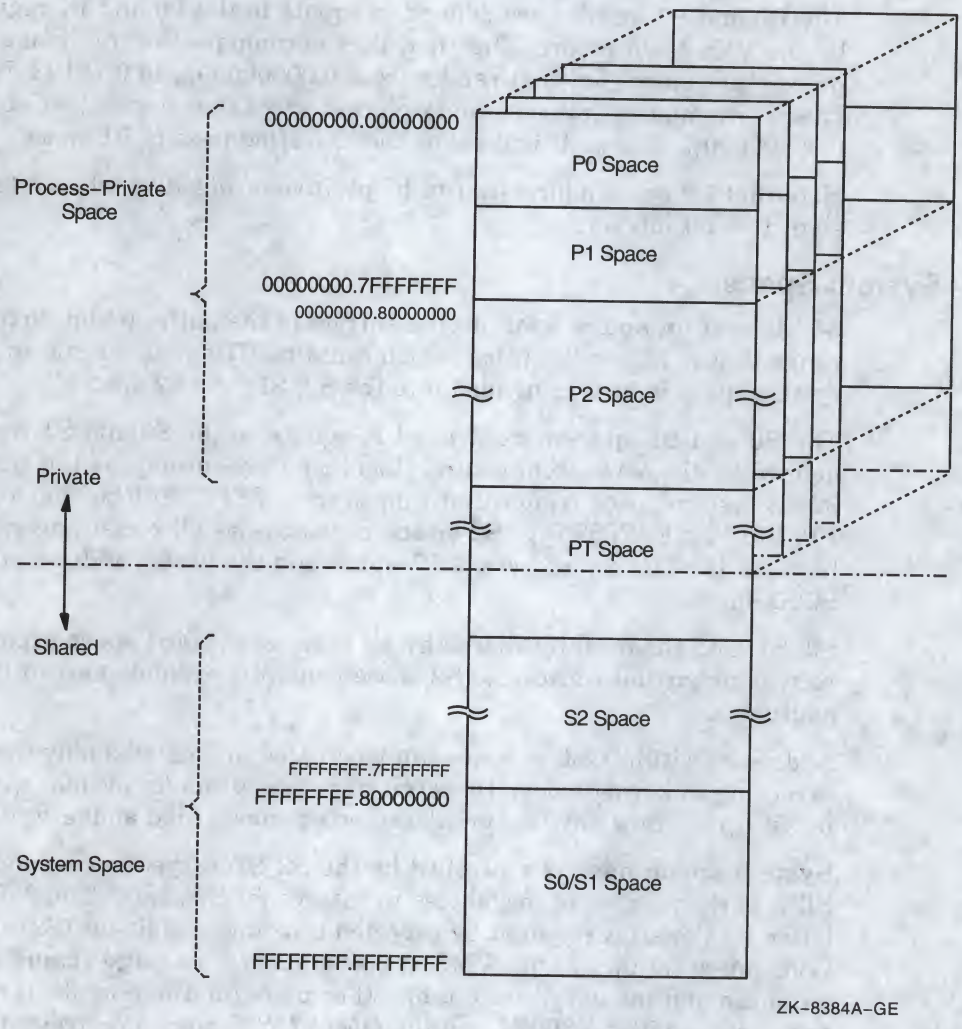

The 64-bit virtual address space layout is designed to accommodate the current and future needs of the OpenVMS Alpha operating system and its users. The new address space consists of the following fundamental areas:

- Process-private space
- System space
- Page table space

1.2.1 Process-Private Space

Supporting process-private address space is a focus of much of the memory management design within the OpenVMS operating system.

Process-private space, or process space, contains all virtual addresses below PT space. As shown in Figure 1–2, the layout of process space is further divided into the P0, P1, and P2 spaces. P0 space refers to the program region. P1 space refers to the control region. P2 space refers to the 64-bit program region.

The **P0** and **P1 spaces** are defined to equate to the P0 and P1 regions defined by the VAX Architecture. Together, they encompass the traditional 32-bit process-private region that ranges from 0.00000000_{16} to $0.7FFFFFFF_{16}$. **P2 space** encompasses all remaining process space that begins just above P1 space, 0.80000000_{16}, and ends just below the lowest address of PT space.

Note that P2 space addresses can be positive or negative when interpreted as signed 64-bit integers.

1.2.2 System Space

64-bit system space refers to the portion of the entire 64-bit virtual address range that is higher than that which contains PT space. As shown in Figure 1–2, system space is further divided into the S0, S1, and S2 spaces.

The **S0** and **S1 spaces** are defined to equate to the S0 and S1 regions defined by the VAX Architecture. Together they encompass the traditional 32-bit system space region that ranges from $FFFFFFFF.80000000_{16}$ to $FFFFFFFF.FFFFFFFF_{16}$. **S2 space** encompasses all remaining system spaces between the highest address of PT space and the lowest address of the combined S0/S1 space.

S0, S1, and S2 are fully shared by all processes. S0/S1 space expands toward increasing virtual addresses. S2 space generally expands toward lower virtual addresses.

Addresses within system space can be created and deleted only from code that is executing in kernel mode. However, page protection for system space pages can be set up to allow any less privileged access mode read and/or write access.

System space base is controlled by the S2_SIZE system parameter. S2_SIZE is the number of megabytes to reserve for S2 space. The default value is based on the sizes required by expected consumers of 64-bit (S2) system space. Consumers set up by OpenVMS at boot time are the **page frame number** (PFN) database and the global page table. (For more information about setting system parameters with SYSGEN, see the *OpenVMS System Management Utilities Reference Manual: M–Z*.)

The **global page table**, also known as the **GPT**, and the PFN database reside in the lowest-addressed portion of S2 space. By moving the GPT and PFN database to S2 space, the size of these areas is no longer constrained to a small portion of S0/S1 space. This allows OpenVMS to support much larger physical memories and much larger global sections.

1.2.3 Page Table Space

In previous versions of the OpenVMS Alpha operating system, **page table space**, also known as **PT space**, was addressable in more than one way. The PALcode TB miss handler used addresses starting at 2.00000000_{16} to read PTEs, while memory management code addressed the page tables primarily within the traditional 32-bit system space. The process page tables were within the process header (PHD), and the system space page tables were located in the highest virtual addresses, all within the traditional 32-bit system space.

As of OpenVMS Alpha Version 7.0, page tables are addressed primarily within 64-bit PT space. Page table references are to this virtual address range; they are no longer in 32-bit shared system address space.

The dotted line in Figure 1–2 marks the boundary between process-private space and shared space. This boundary is in PT space and further serves as the boundary between the process-private page table entries and the shared page table entries. Together, these sets of entries map the entire address space available to a given process. PT space is mapped to the same virtual address for each process, typically a very high address such as $FFFFFFFC.00000000_{16}$.

1.2.4 Virtual Address Space Size

The Alpha architecture supports 64-bit addresses. OpenVMS Alpha Version 7.0 dramatically increases the total amount of virtual address space from 4 GB (gigabytes) to the 8 TB (terabytes) supported by the current Alpha architecture implementations.

The Alpha architecture requires that all implementations must use or check all 64-bits of a virtual address during the translation of a virtual address into a physical memory address. However, implementations of the Alpha architecture are allowed to materialize a subset of the virtual address space. Current Alpha hardware implementations support 43 significant bits within a 64-bit virtual address. This results in an 8 TB address space.

On current Alpha architecture implementations, bit 42 within a virtual address must be sign-extended or propagated through bit 63 (the least significant bit is numbered from 0). Virtual addresses where bits 42 through 63 are not all zeros or all ones result in an access violation when referenced. Therefore, the valid 8 TB address space is partitioned into two disjoint 4 TB ranges separated by a "no access" range in the middle.

The layout of the OpenVMS Alpha address space transparently places this no access range within P2 space. (The OpenVMS Alpha memory management system services always return virtually contiguous address ranges.) The result of the OpenVMS Alpha address space layout design is that valid addresses in P2 space can be positive or negative values when interpreted as signed 64-bit integers.

Note that to preserve 32-bit nonprivileged code compatibility, bit 31 in a valid 32-bit virtual address can still be used to distinguish an address in P0/P1 space from an address in S0/S1 space.

1.3 Virtual Regions

A **virtual region** is a reserved range of process-private virtual addresses. It may be either a "user-defined" region reserved by the user program at run time or a "process-permanent" region reserved by the system on behalf of the process during process creation.

Three process-permanent regions are defined by OpenVMS at the time the process is created:

- Program region (in P0 space)
- Control region (in P1 space)
- 64-bit program region (in P2 space)

These three process-permanent regions exist so that programmers do not have to create regions if their application does not need to reserve additional ranges of address space.

Virtual regions promote modularity within applications by allowing different components of the application to manipulate data in different virtual regions. When a virtual region is created, the caller of the service is returned a region ID to identify that region. The region ID is used when creating, manipulating, and deleting virtual addresses within that region. Different components within an application can create separate virtual regions so that their use of virtual memory does not conflict.

Virtual regions exhibit the following characteristics.

- A virtual region is a "light-weight" object. That is, it does not consume pagefile quota or working set quota for the virtual addresses specified. Creating a user-defined region by calling a new OpenVMS system service merely defines a virtual address range as a distinct address object within which address space can be created, manipulated, and deleted.
- Virtual regions do not overlap. When creating address space within a region, the programmer must specify a region ID to the OpenVMS system service. The programmer must be explicit in which region the address space is to be created.
- The programmer cannot create, manipulate, or delete address space that does not lie entirely within the bounds of a defined region.
- Each user-defined region's size is fixed at the time it is created. Given the large range of virtual addresses in P2 space and the light-weight nature of regions, it is not costly to reserve more address space than the application component immediately needs within that region.

  Note the exception of process-permanent regions, which have no fixed size. See Section 1.3.2.
- Each virtual region has an owner mode and a create mode associated with it. Access modes that are less privileged than the owner of the region cannot delete the region. Access modes that are less privileged than the create mode set for the region cannot create virtual addresses within the region. Owner and create modes are set at the time the region is created and cannot be changed. The create mode for a region cannot be more privileged than the owner mode.

- When virtual address space is created within a region, allocation generally occurs within the region in a densely expanding manner, as is done within the program (P0 space) and control (P1 space) regions. At the time it is created, each region is set up for the virtual addresses within that region to expand toward either increasing virtual addresses, like P0 space, or decreasing virtual addresses, like P1 space. Users can override this allocation algorithm by explicitly specifying starting addresses.
- All user-defined regions are deleted along with the pages created within each region at image rundown.

1.3.1 Regions Within P0 Space and P1 Space

There is one process-permanent region for all of P0 space that starts at virtual address 0 and ends at virtual address $0.3FFFFFFF_{16}$. This is called the **program region**. There is also one process-permanent region for all of P1 space that starts at virtual address 0.40000000_{16} and ends at virtual address $0.7FFFFFFF_{16}$. This is called the **control region**.

The program and control regions are considered to be owned by kernel mode and have a create mode of user, since user mode callers can create virtual address space within these regions. This is upwardly compatible with previous releases of OpenVMS.

These program and control regions cannot be deleted. They are considered to be **process-permanent**.

1.3.2 64-Bit Program Region

P2 space has a densely expandable region starting at the lowest virtual address of P2 space, 0.80000000_{16}. This region is called the **64-bit program region**. Having a 64-bit program region in P2 space allows an application that does not need to take advantage of explicit virtual regions to avoid incurring the overhead of creating a region in P2 space. This region always exists, so addresses can be created within P2 space immediately.

The 64-bit program region is the only region whose size is not fixed when it is created. At process creation, the 64-bit program region encompasses all of P2 space. When a user-defined region is created in P2 space, OpenVMS memory management shrinks the 64-bit program region so that no two regions overlap. When a user-defined region is deleted, the 64-bit program region expands to encompass the virtual addresses within the deleted region if no other user-defined region exists at lower virtual addresses.

As described in Section 1.3.3, a user can create a region in otherwise unoccupied P2 space. If the user-defined region is specified to start at the lowest address of the 64-bit program region, then any subsequent attempt to allocate virtual memory within the region will fail.

The region has a user create mode associated with it, that is, any access mode can create virtual address space within it.

The 64-bit program region cannot be deleted. It is considered to be process-permanent and survives image rundown. Note that all created address space within the 64-bit program region is deleted and the region is reset to encompass all of P2 space as a result of image rundown.

1.3.3 User-Defined Virtual Regions

A user-defined virtual region is a region created by calling the new OpenVMS SYS$CREATE_REGION_64 system service. The location at which a user-defined region is created is generally unpredictable. In order to maximize the expansion room for the 64-bit program region, OpenVMS memory management allocates virtual regions starting at the highest available virtual address in P2 space that is lower than any existing user-defined region.

For maximum control of the process-private address space, the application programmer can specify the starting virtual address when creating a virtual region. This is useful in situations when it is imperative that the user be able to specify exact virtual memory layout.

Virtual regions can be created so that allocation occurs with either increasing addresses or decreasing virtual addresses. This allows applications with stacklike structures to create virtual address space and expand naturally.

Virtual region creation gives OpenVMS subsystems and the application programmer the ability to reserve virtual address space for expansion. For example, an application can create a large virtual region and then create some virtual addresses within that region. Later, when the application requires more virtual address space, it can expand within the region and create more address space in a virtually contiguous manner to the previous addresses allocated within that region.

If you do not explicitly delete a virtual region with the SYS$DELETE_REGION_64 system service, the user-defined region along with all created address space is deleted when the image exits.

1.4 Using 64-Bit Addresses in Applications

Nonprivileged programs can optionally be modified to take advantage of 64-bit addressing features. OpenVMS Alpha 64-bit virtual addressing does not affect nonprivileged programs that are not explicitly modified to exploit 64-bit support. Binary and source compatibility of existing nonprivileged programs is guaranteed.

By using 64-bit addressing capabilities, application programs can map large amounts of data into memory to provide high levels of performance and make use of very large memory (VLM) systems. In addition, 64-bit addressing allows for more efficient use of system resources, allowing for larger user processes as well as higher numbers of users and client/server processes for virtually unlimited scalability.

The remaining chapters in this guide describe the OpenVMS Alpha programming features that you can use to enhance applications to support 64-bit addressing.

1.5 Language and Pointer Support for 64-Bit Addressing

Full support in DEC C and the DEC C Run-Time Library (RTL) for 64-bit addressing make C the preferred language for programming 64-bit applications, libraries, and system code for OpenVMS Alpha. The 64-bit pointers can be seamlessly integrated into existing C code, and new 64-bit applications can be developed, with natural C coding styles, that take advantage of the 64-bit address space provided by OpenVMS Alpha.

Support for all 32-bit pointer sizes (the default), all 64-bit pointer sizes, and the mixed 32-bit and 64-bit pointer size environment provide compatibility as well as flexibility for programming 64-bit OpenVMS applications in DEC C.

1.5 Language and Pointer Support for 64-Bit Addressing

The ANSI compliant "#pragma" approach for supporting the mixed 32-bit and 64-bit pointer environment is common to Digital UNIX. Features of 64-bit C support include memory allocation routine name mapping (transparent support for _malloc64 and _malloc32) and C type checking for 32-bit versus 64-bit pointer types.

The *OpenVMS Calling Standard* describes the techniques used by all OpenVMS languages for invoking routines and passing data between them. The standard also defines the mechanisms that ensure consistency in error and exception handling routines.

The *OpenVMS Calling Standard* has always specified 64-bit wide parameters. In earlier releases of OpenVMS Alpha, called routines almost always ignore the upper 32-bits of arguments. As of OpenVMS Alpha Version 7.0, the *OpenVMS Calling Standard* provides the following support for 64-bit addresses:

- Called routines can start to use complete 64-bit addresses.
- Callers can pass either 32-bit or 64-bit pointers.
- Pointers passed by reference often require a new 64-bit variant of the original routine.
- "Self-identifying" structures, such as those defined for descriptors and item lists, enable an existing API to be enhanced compatibly.

OpenVMS Alpha 64-bit addressing support for mixed pointers also includes the following features:

- OpenVMS Alpha 64-bit virtual address space layout that applies to all processes. (There are no special 64-bit processes or 32-bit processes.)
- 64-bit pointer support for addressing the entire 64-bit OpenVMS Alpha address space layout including P0, P1, and P2 address spaces and S0/S1, S2, and page table address spaces.
- 32-bit pointer compatibility for addressing P0, P1, and S0/S1 address spaces.
- Many new 64-bit system services which support P0, P1, and P2 space addresses.
- Many existing system services enhanced to support 64-bit addressing.
- 32-bit sign-extension checking for all arguments passed to 32-bit pointer only system serivces.
- C and MACRO-32 macros for handling 64-bit addresses.

2

System Services Support for 64-Bit Addressing

New OpenVMS system services are available, and many existing services have been enhanced to manage 64-bit address space. This chapter describes the OpenVMS Alpha system services that support 64-bit addressing. It explains the changes made to 32-bit services to support 64-bit addresses, and it lists the new 64-bit system services.

To see examples of system services that support 64-bit addressing in an application program, see Appendix B. For complete information about the OpenVMS system services listed in this chapter, see the *OpenVMS System Services Reference Manual: A–GETMSG* and *OpenVMS System Services Reference Manual: GETQUI–Z*.

2.1 System Services Terminology

The following system services definitions are used throughout this guide.

32-bit system service
A 32-bit system service is a system service that only supports 32-bit addresses on any of its arguments that specify addresses. If passed by value, on OpenVMS Alpha a 32-bit virtual address is actually a 64-bit address that is sign-extended from 32-bits.

64-bit friendly interface
A 64-bit friendly interface is an interface that can be called with all 64-bit addresses. A 32-bit system service interface is 64-bit friendly if, without a change in the interface, it needs no modification to handle 64-bit addresses. The internal code that implements the system service might need modification, but the system service interface will not.

The majority of OpenVMS Alpha system services prior to OpenVMS Alpha Version 7.0 have 64-bit friendly interfaces for the following reasons:

- The *OpenVMS Calling Standard* defines arguments to standard routines to be 64 bits wide. The caller of a routine sign-extends 32-bit arguments to be 64 bits.
- 64-bit string descriptors can be distinguished from 32-bit string descriptors at run time. (See the *OpenVMS Calling Standard* for more information about 64-bit descriptors.)
- User visible RMS data structures contain embeded type information such that the RMS routines can tell whether an extended form of a structure is being used. (See Chapter 3 for more details about RMS 64-bit addressing support.)

Examples of 64-bit friendly system services are $QIO, $SYNCH, $ENQ, and $FAO.

Examples of routines with 64-bit unfriendly interfaces are most of the memory management system services, such as $CRETVA, $DELTVA, and $CRMPSC. The INADR and RETADR argument arrays do not promote easily to hold 64-bit addresses.

64-bit system service
A 64-bit system service is a system service that is defined to accept all address arguments as 64-bit addresses (not necessarily 32-bit sign-extended values). Also, a 64-bit system service uses the entire 64 bits of all virtual addresses passed to it.

The 64-bit system services include the **_64** suffix for services that accept 64-bit addresses by reference. For promoted services, this distinguishes the 64-bit capable version from its 32-bit counterpart. For new services, it is a visible reminder that a 64-bit wide address cell will be read/written. This is also used when a structure is passed which contains an embedded 64-bit address, if the structure is not self-identifying as a 64-bit structure. Hence, a routine name need not include "_64" simply because it receives a 64-bit decriptor. Remember that passing an arbitrary value by reference does not mean the suffix is required; passing a 64-bit address by reference does.

2.2 New and Enhanced 64-Bit System Services

Table 2–1 summarizes the OpenVMS Alpha system services that support 64-bit addresses. It includes system services from previous releases that have been enhanced to handle 64-bit addresses as well as new OpenVMS Alpha 64-bit system services.

Although RMS system services provide some 64-bit addressing capabilities, they are not listed in this table because they are not full 64-bit system services. See Chapter 3 for more details.

Table 2–1 64-Bit System Services

Service	Arguments
Alignment System Services	
$GET_ALIGN_FAULT_DATA	buffer_64, buffer_size, return_size_64
$GET_SYS_ALIGN_FAULT_DATA	buffer_64, buffer_size, return_size_64
$INIT_SYS_ALIGN_FAULT_REPORT	match_table_64, buffer_size, flags
AST System Service	
$DCLAST	astadr_64, astprm_64, acmode

(continued on next page)

Table 2–1 (Cont.) 64-Bit System Services

Service	Arguments
Condition Handling System Services	
$FAO	ctrstr_64, outlen_64, outbuf_64, p1_64...pn_64
$FAOL	ctrstr_64, outlen_64, outbuf_64, long_prmlst_64
$FAOL_64	ctrstr_64, outlen_64, outbuf_64, quad_prmlst_64
$GETMSG	msgid, msglen_64, bufadr_64, flags, outadr_64
$PUTMSG	msgvec_64, actrtn_64, facnam_64, actprm_64
$SIGNAL_ARRAY_64	mcharg, sigarg_64
Event Flag System Service	
$READEF	efn, state_64
Fast-I/O System Services	
$IO_CLEANUP	fandle
$IO_PERFORM	fandle, chan, iosadr, bufadr, buflen, porint
$IO_PERFORMW	fandle, chan, iosadr, bufadr, buflen, porint
$IO_SETUP	func, bufobj, iosobj, astadr, flags, return_fandle
I/O System Services	
$QIO(W)[1]	efn, chan, func, iosb_64, astadr_64, astprm_64, p1_64, p2_64, p3_64, p4_64, p5_64, p6_64
$SYNCH	efn, iosb_64
Locking System Services	
$DEQ	lkid, vablk_64, acmode, flags
$ENQ(W)	efn, lkmode, lksb_64, flags, resnam_64, parid, astadr_64, astprm_64, blkast_64, acmode
Memory Management System Services	
$CREATE_REGION_64	length_64, region_prot, flags, return_region_id_64, return_va_64, return_length_64, ...
$DELETE_REGION_64	region_id_64, acmode, return_va_64, return_length_64
$GET_REGION_INFO	function_code, region_id_64, start_va_64, ,buffer_length, buffer_address_64, return_length_64
$EXPREG_64	region_id_64, length_64, acmode, flags, return_va_64, return_length_64
$CRETVA_64	region_id_64, start_va_64, length_64, acmode, flags, return_va_64, return_length_64
$CRMPSC_FILE_64	region_id_64, file_offset_64, length_64, chan, acmode, flags, return_va_64, return_length_64, ...

[1]For more information about the $QIO(W) arguments that support 64-bit addressing, see Chapter 5.

(continued on next page)

Table 2–1 (Cont.) 64-Bit System Services

Service	Arguments
Memory Management System Services	
$CRMPSC_PFN_64	region_id_64, start_pfn, page_count, acmode, flags, return_va_64, return_length_64, ...
$UPDSEC_64(W)	start_va_64, length_64, acmode, updflg, efn, iosa_64, return_va_64, return_length_64, ...
$DELTVA_64	region_id_64, start_va_64, length_64, acmode, return_va_64, return_length_64
$CREATE_GFILE	gsdnam_64, ident_64, file_offset_64, length_64, chan, acmode, flags, return_length_64, ...
$CREATE_GPFILE	gsdnam_64, ident_64, prot, length_64, acmode, flags
$CREATE_GPFN	gsdnam_64, ident_64, prot, start_pfn, page_count, acmode, flags
$DGBLSC	flags, gsdnam_64, ident_64
$MGBLSC_64	gsdnam_64, ident_64, region_id_64, section_offset_64, length_64, acmode, flags, return_va_64, return_length_64, ...
$MGBLSC_GPFN_64	gsdnam_64, ident_64, region_id_64, relative_page, page_count, acmode, flags, return_va_64, return_length_64, ...
$CRMPSC_GFILE_64	gsdnam_64, ident_64, file_offset_64, length_64, chan, region_id_64, section_offset, acmode, flags, return_va_64, return_length_64, ...
$CRMPSC_GPFILE_64	gsdnam_64, ident_64, prot, length_64, region_id_64, section_offset_64, acmode, flags, return_va_64, return_length_64, ...
$CRMPSC_GPFN_64	gsdnam_64, ident_64, prot, start_pfn, page_count, region_id_64, relative_page, acmode, flags, return_va_64, return_length_64, ...
$ADJWSL	pagcnt, wsetlm_64
$LKWSET_64	start_va_64, length_64, acmode, return_va_64, return_length_64
$ULWSET_64	start_va_64, length_64, acmode, return_va_64, return_length_64
$PURGE_WS	start_va_64, length_64
$LCKPAG_64	start_va_64, length_64, acmode, return_va_64, return_length_64
$ULKPAG_64	start_va_64, length_64, acmode, return_va_64, return_length_64
$CREATE_BUFOBJ_64	start_va_64, length_64, acmode, flags, return_va_64, return_length_64, return_buffer_handle_64
$DELETE_BUFOBJ	buffer_handle_64
$SETPRT_64	start_va_64, length_64, acmode, prot, return_va_64, return_length_64, return_prot_64
CPU Scheduling System Services	
$CPU_CAPABILITIES	cpu_id, select_mask, modify_mask, prev_mask, flags
$PROCESS_AFFINITY	pidadr, prcnam, select_mask, modify_mask, prev_mask, flags
$PROCESS_CAPABILITIES	pidadr, prcnam, select_mask, modify_mask, prev_mask, flags
SET_IMPLICIT_AFFINITY	pidadr, prcnam, state, cpu_id, prev_mask

(continued on next page)

Table 2–1 (Cont.) 64-Bit System Services

Service	Arguments
Time System Services	
$CANTIM	reqidt_64, acmode
$GETTIM	timadr_64
$SETIMR	efn, daytim_64, astadr_64, reqidt_64, flags
Other System Services	
$CMEXEC_64	routine_64, quad_arglst_64
$CMKRNL_64	routine_64, quad_arglst_64

2.3 Sign-Extension Checking

OpenVMS system services not listed in Table 2–1 and all user-written system services that are not explicitly enhanced to accept 64-bit addresses will receive sign-extension checking. Any argument passed to these services that is not properly sign-extended will cause the error status SS$_ARG_GTR_32_BITS to be returned.

2.4 Language Support for 64-Bit System Services

C function prototypes for system services are available in SYS$LIBRARY:SYS$STARLET_C.TLB (or STARLET). To take advantage of these system service function prototypes, you must explicitly enable them. For more information about using the system service C function prototypes, see the *OpenVMS Version 7.0 New Features Manual*. For more information about C programming support for 64-bit addressing, see Chapter 8.

No new 64-bit MACRO-32 macros are available for system services. The MACRO-32 caller must use the new AMACRO built-in EVAX_CALLG_64 on the new $CALL64 macro. For more information about MACRO-32 programming support for 64-bit addressing, see Chapter 9.

Table 2–1 (Cont.) New 64-Bit System Services

Service	Description
Time-Saving Services	
[illegible]	[illegible]
[illegible]	[illegible]
[illegible]	[illegible]
Miscellaneous Services	
[illegible]	[illegible]
[illegible]	[illegible]

2.3 Sign-Extension Checking

[illegible]

2.4 Language Support for 64-Bit System Services

[illegible]

3

RMS Interface Enhancements for 64-Bit Addressing

This chapter summarizes changes to the RMS interface that support 64-bit addressing and enable you to use RMS to perform input and output operations to P2 or S2 space. You can take full advantage of these RMS enhancements by making only minor modifications to existing RMS code.

For complete information about RMS support for 64-bit addressing, see the *OpenVMS Record Management Services Reference Manual*.

The RMS user interface consists of a number of control data structures (FAB, RAB, NAM, XABs). These are linked together with 32-bit pointers and contain embedded pointers to I/O buffers and various user data buffers, including file name strings and item lists. RMS support for 64-bit addressable regions allows 64-bit addresses for the following user I/O buffers:

- UBF (user record buffer)
- RBF (record buffer)
- RHB (fixed-length record header buffer; fixed portion of VFC record format)
- KBF (key buffer containing the key value for random access)

The prompt buffer pointed to by RAB$L_PBF is excluded because the terminal driver does not allow 64-bit addresses.

Specific enhancements to the RMS interface for 64-bit addressing are as follows:

- Data buffers can be placed in P2 or S2 space for the following user I/O services:
  - Record I/O services: $GET, $FIND, $PUT, $UPDATE
  - Block I/O services: $READ, $WRITE
- The RAB structure points to the record and data buffers used by these services.
- An extension of the existing RAB structure is used to specify 64-bit buffer pointers and sizes.
- The buffer size maximum for RMS block I/O services ($READ and $WRITE) has been increased from 64 KB bytes to 2 GB bytes, with two exceptions:
  - For RMS journaling, a journaled $WRITE service is restricted to the current maximum (65535 minus 99 bytes of journaling overhead). An RSZ error is returned to RAB$L_STS if the maximum is exceeded.
  - Magnetic tape is still limited to 65535 bytes at the device driver level.

  Buffer size maximums for RMS record I/O services ($GET, $PUT, $UPDATE) have not changed. Prior RMS on-disk record size limits still apply.

The rest of the RMS interface currently is restricted to 32-bit pointers:

- FAB, RAB, NAM, and XABs must still be allocated in 32-bit space.
- Any descriptors or embedded pointers to file names, item lists, and so on, must continue to use 32-bit pointers.
- Any arguments passed to the RMS system services remain 32-bit arguments. If you attempt to pass a 64-bit argument, the SS$_ARG_GTR_32_BITS error is returned.

3.1 RAB64 Data Structure

The RAB64, a new RMS user interface structure, is an extended RAB that can accommodate 64-bit buffer addresses. The RAB64 data structure consists of a 32-bit RAB structure followed by a 64-bit extension as shown below:

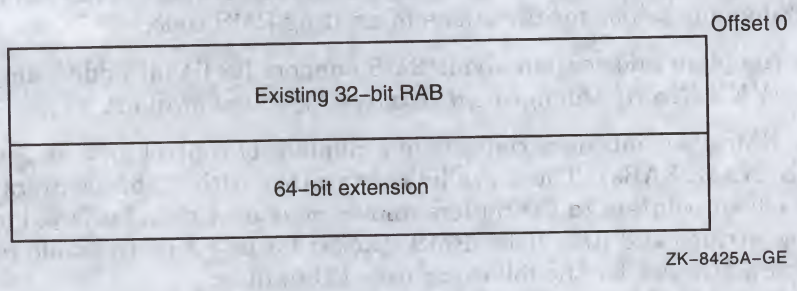

The RAB64 contains fields identical to all of the RAB fields except that field names have the RAB64 prefix instead of the RAB prefix. In addition, RAB64 has the following new fields in the extension:

This new field...	...is an extension of this field	Description
RAB64$Q_CTX	RAB64$L_CTX	User context. This field is not used by RMS but is available to the user. The CTX field is often used to contain a pointer. For asynchronous I/O, it provides the user with the equivalent of an AST parameter.
RAB64$PQ_KBF	RAB64$L_KBF	Key buffer address containing the key value for random access (for $GET and $FIND).
RAB64$PQ_RBF	RAB64$L_RBF	Record buffer address (for $PUT, $UPDATE, and $WRITE).
RAB64$PQ_RHB	RAB64$L_RHB	Record header buffer address (fixed portion of VFC record format).
RAB64$Q_RSZ	RAB64$W_RSZ	Record buffer size.
RAB64$PQ_UBF	RAB64$L_UBF	User buffer address (for $GET and $READ).
RAB64$Q_USZ	RAB64$W_USZ	User buffer size.

Note that the fields with the PQ tag in their names can hold either a 64-bit address or a 32-bit address sign-extended to 64 bits. Therefore, you can use the new fields in all applications whether or not you are using 64-bit addresses.

For most record I/O service requests, there is an RMS internal buffer between the device and the user's data buffer. The one exception occurs with the RMS service $PUT. If the device is a unit record device and it is not being accessed over the network, RMS passes the address of the user record buffer (RBF) to the $QIO system service. If you inappropriately specify a record buffer (RBF)

allocated in 64-bit address space for a $PUT to a unit record device that does not support 64-bit address space (for example, a terminal), the $QIO service returns SS$_NOT64DEVFUNC. (See Chapter 5 for more information about $QIO.) RMS returns the error status RMS$_SYS with SS$_NOT64DEVFUNC as the secondary status value in RAB64$L_STV.

RMS system services support the RAB structure as well as the new RAB64 structure.

3.2 Using the 64-Bit RAB Extension

Only minimal source code changes are required for applications to use 64-bit RMS support.

RMS allows you to use the RAB64 wherever you can use a RAB. For example, a RAB64 can be used in place of a RAB as the first argument passed to any of the RMS record or block I/O services.

Because the RAB64 is an upwardly compatible extension of the existing RAB, most source modules can treat references to fields in a RAB64 as if they were references to a RAB. The 64-bit buffer address counterpart is used only if the following two conditions are met:

- The RAB64$B_BLN field has been initialized to RAB64$C_BLN64 to show that the extension is present.
- The 32-bit address field in the 32-bit portion of the RAB contains −1.

The value in the new quadword size field is used only if the contents of the 32-bit address field designate its use. For example:

If this address field contains -1	The address in this field is used	And the size in this field is used
RAB64$L_UBF	RAB64$PQ_UBF[1]	RAB64$Q_USZ
RAB64$L_RBF	RAB64$PQ_RBF[1]	RAB64$Q_RSZ
RAB64$L_KBF	RAB64$PQ_KBF	RAB64$B_KSZ
RAB64$L_RHB	RAB64$PQ_RHB	FAB$B_FSZ

[1]This field can contain either a 64-bit address or a 32-bit address sign-extended to 64 bits.

While RMS allows you to use the RAB64 wherever you can use a RAB, some source languages may impose other restrictions. Consult the documentation for your source language for more information.

3.3 New Macros

The following new MACRO-32 and BLISS macros have been implemented to support the 64-bit extension to the user RAB structure:

- MACRO-32 macros
  - $RAB64 (counterpart to $RAB)
  - $RAB64_STORE (counterpart to $RAB_STORE)

  Using these macros has the following results:
  - RAB$B_BLN is assigned the constant of RAB$C_BLN64.

- The original longword I/O buffers are initialized to −1, and the USZ and RSZ word sizes are initialized to 0.
- Values specified using the UBF, USZ, RBF, RSZ, RHB, or KBF keywords are moved into the quadword fields for these keywords. (In contrast, the $RAB and $RAB_STORE macros move these values into the longword [or word] fields for these keywords.)

- BLISS macros

  The following BLISS macros are available only in the STARLET.R64 library because they use the QUAD keyword, which is available only to BLISS-64. Thus, any BLISS routines referencing them must be compiled using the BLISS-64 compiler.

  - $RAB64 (counterpart to $RAB)
  - $RAB64_INIT (counterpart to $RAB_INIT)
  - $RAB64_DECL (counterpart to $RAB_DECL)

  Using the first two macros has these results:

  - RAB$B_BLN is assigned the constant of RAB$C_BLN64.
  - The original longword I/O buffers are initialized to −1, and the USZ and RSZ word sizes are initialized to 0.
  - Values assigned to the keywords UBF, USZ, RBF, RSZ, RHB, or KBF are moved into the quadword fields for these keywords. (In contrast, the $RAB and $RAB_INIT macros move these values into the longword [or word] fields for these keywords.)

  The third macro allocates a block structure of bytes with a length of RAB$C_BLN64.

4

File System Support for 64-Bit Addressing

The Extended QIO Processor (XQP) file system, which implements the Files–11 On-Disk Structure Level 2 (ODS-2) and the Magnetic Tape Ancillary Control Process (MTAAACP) both provide support for the use of 64 bit buffer addresses for virtual read and write functions.

The XQP and ACP translate a virtual I/O request to a file into one or more logical I/O requests to a device. Because the buffer specified with the XQP or ACP request is passed on to the device driver, the support for buffers in P2 or S2 space is also dependent on the device driver used by the XQP and ACP.

All OpenVMS supplied disk and tape drivers support 64-bit addresses for data transfers to and from disk and tape devices on the virtual, logical, and physical read and write functions. Therefore, the XQP and Magnetic Tape ACP support buffers in P2 or S2 space on the virtual read and write functions.

The XQP and ACP do not support buffer addresses in P2 or S2 space on the control functions (IO$_ACCESS, IO$_DELETE, IO$_MODIFY, and so on).

For more information about device drivers that support 64-bit buffer addresses, see Chapter 5.

File System Support for 64-Bit Addressing

5

OpenVMS Alpha Device Support for 64-Bit Addressing

Input and output operations can be performed directly to and from P2 or S2 space by means of RMS services, the $QIO system service, and most of the device drivers supplied with OpenVMS Alpha systems.

This chapter explains how the $QIO system service supports 64-bit addresses, describes the OpenVMS Alpha device drivers that do and do not support 64-bit addresses, and lists the OpenVMS Alpha disk and tape driver function codes that support 64-bit addresses.

Customer-written device drivers can be modified to support 64-bit addresses. For information about how to modify a customer-written device driver to support 64-bit addressing, refer to the *OpenVMS Alpha Guide to Upgrading Privileged-Code Applications*. To see an example of a device driver that has been modified to support a 64-bit buffer address in all of its functions, refer to the LRDRIVER device driver in the SYS$EXAMPLES directory.

Important

OpenVMS Alpha Version 7.0 includes significant changes to OpenVMS Alpha privileged interfaces and data structures. As a result of these changes, all device drivers from previous versions of OpenVMS Alpha (including field test versions of OpenVMS Alpha Version 7.0) must be recompiled and relinked to run correctly on OpenVMS Alpha Version 7.0.

For more details about OpenVMS Alpha Version 7.0 changes that may require source changes to customer-written drivers, see the *OpenVMS Alpha Guide to Upgrading Privileged-Code Applications*.

5.1 $QIO Support for 64-Bit Addresses

The $QIO and $QIOW system services accept the following arguments:

```
$QIO[W] efn,chan,func,iosb,astadr,astprm,p1,p2,p3,p4,p5,p6
```

These services have a 64-bit friendly interface (as described in Chapter 2), which allows these services to support 64-bit addresses.

Table 5–1 summarizes the changes to the data types of the $QIO and $QIOW system service arguments to accommodate 64-bit addresses.

Table 5–1 $QIO[W] Argument Changes

Argument	Prior Type	New Type	Description
efn	Unsigned longword	-	Event flag number. Unchanged.
chan	Unsigned word	-	Channel number. Unchanged.
func	Unsigned longword	-	I/O function code. Unchanged.
iosb	32-bit pointer[1]	64-bit pointer	Pointer to a quadword I/O status block (IOSB). The IOSB format is unchanged.
astadr	32-bit pointer[1]	64-bit pointer	Procedure value of the caller's AST routine. On Alpha systems, the procedure value is a pointer to the procedure descriptor.
astprm	Unsigned longword[2]	Quadword	Argument value for the AST routine.
P1	Longword[2]	Quadword	Device-dependent argument. Often P1 is a buffer address.
P2	Longword[2]	Quadword	Device-dependent argument. Only the low-order 32-bits will be used by system-supplied FDT routines that use P2 as the buffer size.
P3	Longword[2]	Quadword	Device-dependent argument.
P4	Longword[2]	Quadword	Device-dependent argument.
P5	Longword[2]	Quadword	Device-dependent argument.
P6	Longword[2]	Quadword	Device-dependent argument. Sometimes P6 is used to contain the address of a diagnostic buffer.

[1]32-bit pointer was sign-extended to 64 bits as required by the *OpenVMS Calling Standard*.

[2]32-bit longword value was sign-extended to 64 bits as required by the *OpenVMS Calling Standard*.

Usually the $QIO P1 argument specifies a buffer address. All the system-supplied upper-level FDT routines that support the read and write functions use this convention. The P1 argument determines whether the caller of the $QIO service requires 64-bit support. If the $QIO system service rejects a 64-bit I/O request, the following fatal system error status is returned:

```
SS$_NOT64DEVFUNC  64-bit address not supported by device for this function
```

This fatal condition value is returned under the following circumstances:

- The caller has specified a 64-bit virtual address in the P1 device dependent argument, but the device driver does not support 64-bit addresses with the requested I/O function.
- The caller has specified a 64-bit address for a diagnostic buffer, but the device driver does not support 64-bit addresses for diagnostic buffers.
- Some device drivers may also return this condition value when 64-bit buffer addresses are passed using the P2 through P6 arguments and the driver does not support a 64-bit address with the requested I/O function.

For more information about the $QIO, $QIOW, and $SYNCH system services, see the *OpenVMS System Services Reference Manual: GETQUI–Z*.

5.2 OpenVMS Drivers Supporting 64-Bit Addresses

A device driver declares support for 64-bit addresses individually by I/O function code. Disk and tape device drivers support 64-bit addresses for data transfers to and from disk and tape devices on the virtual, logical, and physical read and write functions. For example, the OpenVMS SCSI disk class driver, SYS$DKDRIVER, supports 64-bit addresses on the IO$_READVBLK and IO$_WRITEVBLK functions, but not on the IO$_AUDIO function.

OpenVMS Alpha device drivers that support 64-bit adresses include the following:

- All disk and tape drivers
- All port drivers below disk and tape drivers
- LAN drivers
- Mailbox driver
- ISA parallel port driver (LRDRIVER.C)

Table 5–2 lists the OpenVMS Alpha device drivers that support for 64-bit addresses on at least some functions.

Table 5–2 Drivers Supporting 64-Bit Addresses

Driver	Description
SYS$DADDRIVER	Local area disk client disk driver
SYS$DKDRIVER	SCSI disk class driver
SYS$DUDRIVER	DSA disk class driver
SYS$DVDRIVER	Floppy disk for Intel 83077AA
SYS$ECDRIVER	LAN Driver for PMAI
SYS$ERDRIVER	LAN Driver for DE422
SYS$ESDRIVER	LAN Driver for DESUA
SYS$EWDRIVER	TULIP LAN, PCI
SYS$EXDRIVER	DEMNA LAN, XMI
SYS$EZDRIVER	SGEC/INEC/TGEC LAN
SYS$FADRIVER	FDDI for Futurebus

(continued on next page)

Table 5–2 (Cont.) Drivers Supporting 64-Bit Addresses

Driver	Description
SYS$FCDRIVER	DEFZA, DEFTA LAN, TC
SYS$FRDRIVER	DEFEA LAN, EISA
SYS$FXDRIVER	DEMFA LAN, XMI
SYS$GKDRIVER	SCSI generic class driver
SYS$HCDRIVER	OTTO class ATM
SYS$ICDRIVER	TMS380 LAN, TC
SYS$IRDRIVER	TMS380 EISA Token Ring
SYS$LADDRIVER	Local Area Disk
SYS$LASTDRIVER	Local Area System Trans
SYS$LRDRIVER	VL82C106 parallel printer driver
SYS$MADDRIVER	Local area client tape
MBDRIVER	Mailbox driver
SYS$MKDRIVER	SCSI tape class driver
NLDRIVER	Null device driver
SYS$PADRIVER	SHAC CI and DSSI port driver
SYS$PEDRIVER	NI SCS port driver
SYS$PIDRIVER	NCR 53C710 DSSI port
SYS$PKCDRIVER	SCSI NCR 53C94 Port
SYS$PKEDRIVER	NCR 53C810 SCSI port
SYS$PKJDRIVER	ADAPTEC 1742A SCSI port
SYS$PKSDRIVER	SIMport TC-SCSI port
SYS$PKTDRIVER	NCR 53C710 SCSI port
SYS$PKZDRIVER	XZA SCSI Port
SYS$PNDRIVER	NPORT SCS port
SYS$PUDRIVER	CI UDA port driver
SYS$SHDRIVER	Volume shadowing
SYS$TUDRIVER	MSCP/DSA tape class
SYS$WPDRIVER	Watchpoint driver

Table 5–3 lists the OpenVMS Alpha device drivers that do not support 64-bit addresses in OpenVMS Alpha Version 7.0.

Table 5–3 Drivers Restricted to 32-Bit Addresses

Driver	Description
SYS$CTDRIVER	CTERM driver
SYS$FBDRIVER	Terminal fallback driver
SYS$FTDRIVER	Psuedo terminal driver
SYS$FYDRIVER	DUP DSA protocol class driver

(continued on next page)

Table 5–3 (Cont.) Drivers Restricted to 32-Bit Addresses

Driver	Description
SYS$GQADRIVER	QVISION driver
SYS$GTADRIVER	DECwindows TX driver for Flamingo
SYS$GXADRIVER	Flamingo CXTurbo (aka SFB, aka HX) driver
SYS$GYADRIVER	SFB+ aka HX+, aka FFB driver
SYS$GYBDRIVER	Driver for TGA graphics on the PCI bus
SYS$IEDRIVER	DECwindows extension
SYS$IKBDRIVER	DECwindows PCXAL keyboard
SYS$IKDRIVER	DECwindows LKxxx keyboard
SYS$IMBDRIVER	DECwindows PCXAS (PS2) mouse
SYS$IMDRIVER	DECwindows VSxxx mouse
SYS$INDRIVER	DECwindows input driver
SYS$LTDRIVER	LAT terminal driver
NDDRIVER	DECnet Phase IV DLE (MOP support)
NETDRIVER	DECnet Phase IV
SYS$RTTDRIVER	Remote DECnet terminal driver
SYS$SODRIVER	AMD79C30A Audio/ISDN driver
SYS$TTDRIVER	Terminal class driver
DECW$XTDRIVER	X Terminal class driver
SYS$YRDRIVER	Z85C30 SCC terminal port driver
SYS$YSDRIVER	PC87312 terminal port driver

Some notable points about the drivers that are restricted to 32-bit buffer addresses include the following:

- The terminal drivers do not support 64-bit addresses.
- The drivers used by DECwindows Motif software do not support 64-bit addresses.
- The DECnet Phase IV drivers do not support 64-bit addresses.

5.3 Function Codes that Support 64-Bit Addresses

Table 5–4 lists the OpenVMS Alpha I/O function codes that support 64-bit addresses.

Table 5–4 64-Bit Capable I/O Functions

Driver Type	Function Code	64-Bit Addresses
Disk		
	IO$_READLBLK	P1
	IO$_READPBLK	P1
	IO$_READVBLK	P1
	IO$_WRITECHECK	P1
	IO$_WRITELBLK	P1
	IO$_WRITEPBLK	P1
	IO$_WRITEVBLK	P1
Magnetic Tape		
	IO$_READLBLK	P1
	IO$_READPBLK	P1
	IO$_READVBLK	P1
	IO$_WRITELBLK	P1
	IO$_WRITEOF	P1
	IO$_WRITEPBLK	P1
	IO$_WRITEVBLK	P1
Mailbox		
	IO$_READLBLK	P1
	IO$_READPBLK	P1
	IO$_READVBLK	P1
	IO$_WRITELBLK	P1
	IO$_WRITEPBLK	P1
	IO$_WRITEVBLK	P1
Local Area Network (LAN)		
	IO$_READLBLK	P1,P5
	IO$_READPBLK	P1,P5
	IO$_READVBLK	P1,P5
	IO$_WRITELBLK	P1,P4,P5
	IO$_WRITEPBLK	P1,P4,P5
	IO$_WRITEVBLK	P1,P4,P5

5.4 64-Bit IO$_DIAGNOSE Function for SCSI Class Drivers

The $QIO IO$_DIAGNOSE function has been enhanced to support 64-bit addressing for the following SCSI class drivers: GKDRIVER, DKDRIVER, and MKDRIVER. This means that the virtual addresses specified within the S2DGB may now be 64-bit virtual addresses if the user application requests it.

The $QIO IO$_DIAGNOSE arguments are still as follows:

Argument	Use
P1	S2DGB base address
P2	S2DGB length
P3	Reserved, should be zero
P4	Reserved, should be zero
P5	Reserved, should be zero
P6	Reserved, should be zero

The SCSI Diagnose Buffer (S2DGB) defined in STARLET now allows two formats, one for 32-bit addressing and one one for 64-bit addressing. The 32-bit format is identical to the one supported on OpenVMS Alpha Version 6.2. Figure 5–1 shows the 32-bit S2DGB format. Figure 5–2 shows the 64-bit S2DGB format.

Figure 5–1 OpenVMS SCSI-2 Diagnose Buffer (S2DGB) 32-Bit Layout

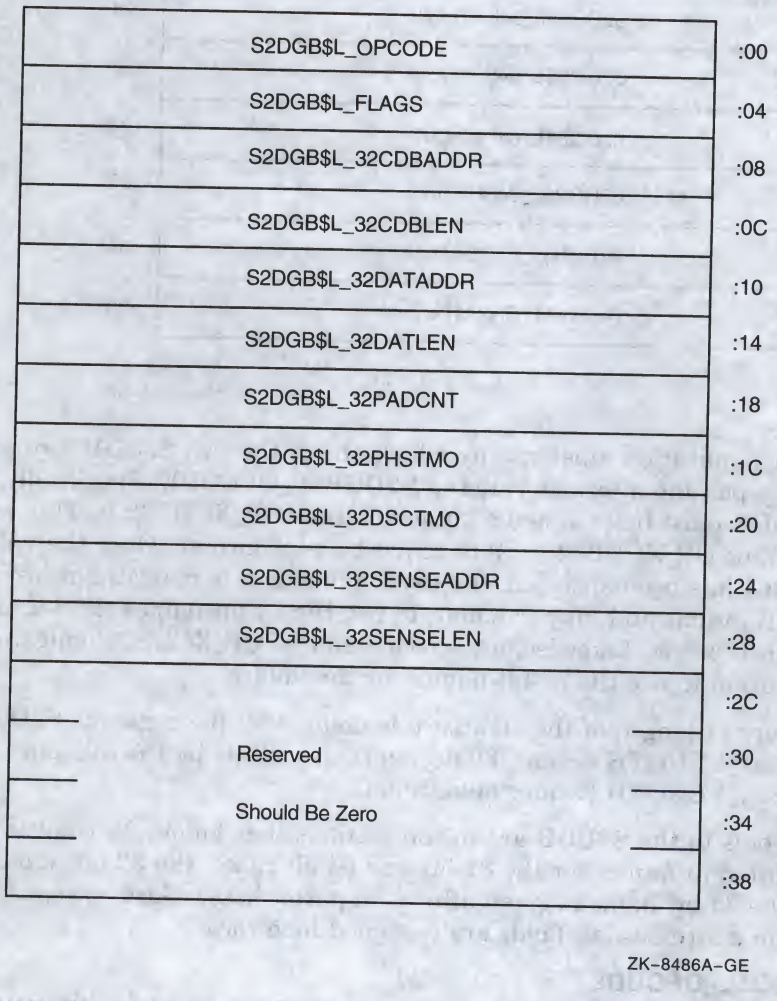

ZK–8486A–GE

Figure 5–2 OpenVMS SCSI-2 Diagnose Buffer (S2DGB) 64-Bit Layout

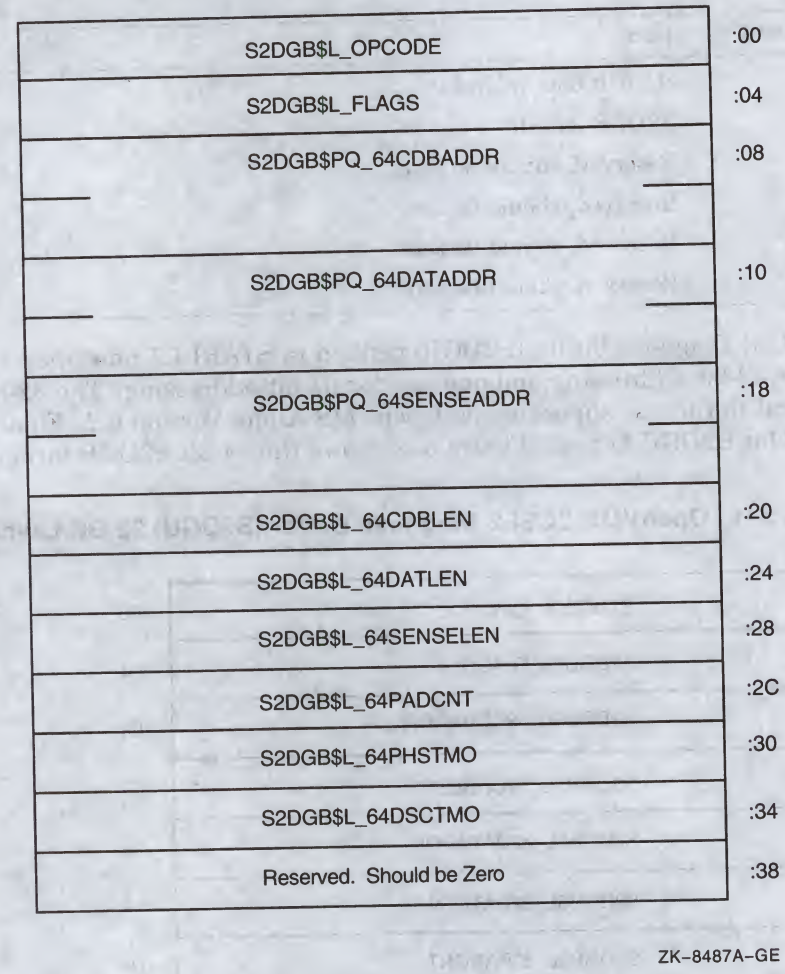

A user application must specify which one of the two S2DGB formats is to be used by passing a format value in S2DGB$L_OPCODE. Specifically, S2DGB$L_OPCODE must be assigned a value of either OP_XCDB32 (= 1) to request 32-bit format, or OP_XCDB64 (= 2) to request 64-bit format. Once the value of OP_XCDB64 has been specified, the user application is obligated to use the 64-bit S2DGB format and, in particular, to use the 64-bit names for S2DGB fields as described below. Likewise, an opcode value of OP_XCDB32 obligates the user application to use the 32-bit names for the fields.

The correct length of the structure is defined by the constant S2DGB$K_XCDB32_LENGTH (value: 60-decimal), as well as by the constant S2DGB$K_XCDB64_LENGTH (value: 60-decimal).

The fields in the S2DGB are in the sections that follow. Whenever a field has two different names for the 32-bit and 64-bit cases, the 32-bit name is given first, and the 64-bit name is given after it in parentheses. Also, except for fields which contain addresses, all fields are unsigned longwords.

S2DGB$L_OPCODE

This field should contain either S2DGB$K_OP_XCDB32 or S2DGB$K_OP_XCDB64, depending on whether the user application intends to supply 32-bit virtual addresses or 64-bit virtual addresses, respectively, in the other fields of the S2DGB.

S2DGB$L_FLAGS

This field should contain the bit fields shown in the following table. Note that these bit definitions start at bit 0 and omit no bits. This is required for compatibility with the IO$_DIAGNOSE interface available in OpenVMS Alpha Version 6.1 and earlier.

S2DGB$V_READ	This bit should be 1 if the operation being performed is a read. If the operation is a write, this bit should be 0.
S2DGB$V_DISCPRIV	This bit should contain the DiscPriv bit value to be used in the IDENTIFY message sent with this operation. If S2DGB$V_TAGGED_REQ is 1, then this bit should be ignored. Note that this bit may be ignored by some ports.
S2DGB$V_SYNCHRONOUS	This bit is ignored since its value is beyond the control of the user in SCSI-2 drivers.
S2DGB$V_OBSOLETE1	This bit is ignored. In previous releases, it represented the disabling of command retries, which is now beyond the control of the user in SCSI-2 drivers.
S2DGB$V_TAGGED_REQ	When this bit is 1, the operation is processed as using tagged command queuing and S2DGB$V_TAG should define the tag value to be used. When this bit is 0, the operation is processed without benefit of tagged command queuing. Ports that do not support tagged command queuing always behave as if this bit is 0. Note that some ports simulate untagged operations using appropriately tagged operations. If S2DGB$V_TAGGED_REQ is 1, then this 3-bit field should contain one of the following coded constant values: S2DGB$K_SIMPLE indicates that the command is to be sent with the SIMPLE queue tag. S2DGB$K_ORDERED indicates that the command is to be sent with the ORDERED queue tag. S2DGB$K_EXPRESS indicates that the command is to be sent with the HEAD OF QUEUE queue tag. If S2DGB$V_TAGGED_REQ is 0, then this field is ignored. Ports that do not support tagged command queuing always ignore the S2DGB$V_TAG field and send all commands as untagged operations. Note that automatic contingent allegiance processing is not accessible through the IO$_DIAGNOSE function. Also, even though this is a 3-bit field, only 2 bits are currently being utilized. That is, the 3 constants above represent values, not bit positions.

S2DGB$V_AUTOSENSE	When this bit is 1, S2DGB$L_32SENSEADDR and S2DGB$L_32SENSELEN should contain a valid sense buffer address and length. If a CHECK CONDITION or COMMAND TERMINATED status is returned, REQUEST SENSE data will be returned in the buffer defined by S2DGB$L_32SENSEADDR and S2DGB$L_32SENSELEN. When S2DGB$V_AUTOSENSE is 0, the buffer described by S2DGB$L_32SENSEADDR and S2DGB$L_32SENSELEN is ignored. In such cases, the class driver saves the autosense data in pool and returns it to the next IO$_DIAGNOSE, if and only if that IO$_DIAGNOSE has a REQUEST SENSE CDB. All other bits in S2DGB$L_FLAGS should be zero.

S2DGB$L_32CDBADDR (S2DGB$PQ_64CDBADDR)

This field should contain the 32-bit (or 64-bit) virtual address of the SCSI command data block (CDB) to be sent to the target by this IO$_DIAGNOSE operation.

Note that S2DGB$L_32CDBADDR is a pointer to a longword, while S2DGB$PQ_64CDBADDR is a pointer to a quadword.

S2DGB$L_32CDBLEN (S2DGB$L_64CDBLEN)

This field should contain the number of bytes in the SCSI command data block (CDB) to be sent to the target by this IO$_DIAGNOSE operation. (Legal values: 2 to 248. However, some ports may restrict CDBs to smaller lengths. Recommended values: 2 to 16.)

S2DGB$L_32DATADDR (S2DGB$PQ_64DATADDR)

This field should contain the 32-bit (or 64-bit) virtual address of the DATAIN or DATAOUT buffer to be used with this SCSI operation. If the CDB being sent to the target does not use a DATAIN or DATAOUT buffer, then this field should be zero.

Note that S2DGB$L_32DATADDR is a pointer to a longword, while S2DGB$PQ_64DATADDR is a pointer to a quadword.

S2DGB$L_32DATLEN (S2DGB$L_64DATLEN)

This field should contain the number of bytes in the DATAIN or DATAOUT buffer associated with this operation. If the CDB being sent to the target does not use a DATAIN or DATAOUT buffer, then this field should be zero. (Legal values: 0 to UCB$L_MAXBCNT. Recommended values: 0 to 65,536. All ports are required to support at least 65,536 byte data transfers.)

S2DGB$L_32PADCNT (S2DGB$L_64PADCNT)

This field should contain the number of padding DATAIN or DATAOUT bytes required by this operation. (Legal values: 0 to the maximum number of bytes in a disk block on this system minus one. Current legal values: 0 to 511.)

S2DGB$L_32PHSTMO (S2DGB$L_64PHSTMO)

This field should contain the number of seconds that the port driver should wait for a phase transition to occur or for delivery of an expected interrupt. If S2DGB$V_ TAGGED_REQ is 1 or this field contains a 0 or 1, then the current phase transition timeout setting will not be changed. (Legal values: 0 to 300 [5 minutes].)

S2DGB$L_32DSCTMO (S2DGB$L_64DSCTMO)
This field should contain the number of seconds that the port driver should wait for a disconnected transaction to reconnect. If S2DGB$V_TAGGED_REQ is 1 or this field contains a 0 or 1, then the current disconnect timeout setting will not be changed. (Legal values: 0 to 65,535 [about 18 hours].)

S2DGB$L_32SENSEADDR (S2DGB$PQ_64SENSEADDR)
If S2DGB$V_AUTOSENSE is 1, then this field should contain the 32-bit (or 64-bit) virtual address of the sense buffer to be used by this SCSI operation. If S2DGB$V_AUTOSENSE is 0, this field will be ignored.

Note that S2DGB$L_32SENSEADDR is a pointer to a longword, while S2DGB$PQ_64SENSEADDR is a pointer to a quadword.

S2DGB$L_32SENSELEN (S2DGB$L_64SENSELEN)
If S2DGB$V_AUTOSENSE is 1, then this field should contain the number of bytes in the sense buffer associated with this operation. (Legal values: 0 to 255. Note: a value of 0 instructs the class driver to discard any sense data received. Recommended value: 18. Some ports may restrict the number of sense bytes to 18.) If S2DGB$V_AUTOSENSE is 0, this field will be ignored.

5.4.1 64-Bit S2DGB Example

The following example shows how to set up a 64-bit S2DGB:

```
#include <s2dgbdef.h>                                          /* Define S2DGB   */
#include <far_pointers.h>                                      /* Define VOID_PQ */

  S2DGB diag_desc;

  /* Set up some default S2DGB descriptor values */

  diag_desc.s2dgb$l_opcode = OP_XCDB64                        /* Use 64-bits  */
  diag_desc.s2dgb$l_flags =  (S2DGB$M_READ |                  /* Flags*/
                      S2DGB$M_TAGGED_REQ |
                      S2DGB$M_AUTOSENSE);
  diag_desc.s2dgb$v_tag = S2DGB$K_SIMPLE;                     /* SIMPLE que tag */
  diag_desc.s2dgb$pq_64cdbaddr = (VOID_PQ)(&cdb[0]);/* Command addr  */
  diag_desc.s2dgb$l_64cdblen =  6;                            /* Command length */
  diag_desc.s2dgb$pq_64dataddr = (VOID_PQ)(&buf[0]);/* Data addr      */
  diag_desc.s2dgb$l_64datlen = 20;                            /* Data length    */
  diag_desc.s2dgb$l_64padcnt = 0;                             /* Pad length    */
  diag_desc.s2dgb$l_64phstmo = 20;                            /* Phase timeout */
  diag_desc.s2dgb$l_64dsctmo = 10;                            /* Disc timeout  */
  diag_desc.s2dgb$pq_64senseaddr = (VOID_PQ)(&asn[0]);/* Autosense addr */
  diag_desc.s2dgb$l_64senselen = 255;                         /* Sense length  */
  diag_desc.s2dgb$l_reserved_1 = 0;                           /* Reserved     */
       .
       .
       .

  status = sys$qiow(0, target_chan, IO$_DIAGNOSE, &iosb, 0, 0,
           &diag_desc,  S2DGB$K_XCDB64_LENGTH,  0,  0,  0,  0);
```

If all arguments are valid, the class driver will invoke the necessary port functions to send the CDB, transfer the data, and return, save or discard sense data as defined by the input S2DGB. Upon completion, the return IOSB will have the following format:

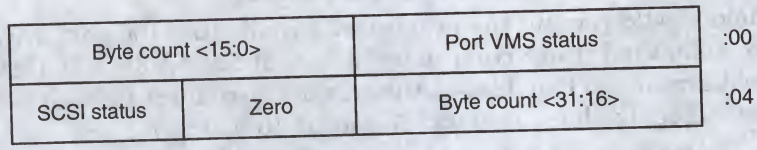

The DKDRIVER, GKDRIVER, and MKDRIVER class drivers, which implement other QIO functions, might intermix other tagged requests with IO$_DIAGNOSE requests. The order in which requests are sent generally matches the order in which requests are presented to the driver. An exception to this ordering occurs when the driver receives REQUEST SENSE for which autosense data previously has been recovered and stored. In this case, the IO$_DIAGNOSE will complete immediately and no command will be sent to the target.

The DKDRIVER, GKDRIVER, and MKDRIVER class drivers permit only one IO$_DIAGNOSE operation to be active (in the start I/O routine) at given time, except as described in the next paragraph. However, applications must single thread IO$_DIAGNOSE requests in order to properly detect the presence of sense data and send the required REQUEST SENSE command. This is consistent with the VAX IO$_DIAGNOSE behavior. For example, if three reads are issued with no waiting and the first read gets a CHECK CONDITION, the sense data will be discarded by the target when the second read arrives.

The DKDRIVER, GKDRIVER, and MKDRIVER drivers permit more than one IO$_DIAGNOSE operation to be active (in the start I/O routine) only when all active operations have the S2DGB$V_AUTOSENSE flag equal to 1. Upon encountering the first IO$_DIAGNOSE with S2DGB$V_AUTOSENSE equal to 0, the class driver will apply the restrictions described in the previous paragraph.

6 OpenVMS Alpha 64-Bit API Guidelines

This chapter describes the guidelines used to develop 64-bit friendly interfaces to support OpenVMS Alpha 64-bit virtual addressing. Application programmers who are developing their own 64-bit application programming interfaces might find this information useful.

These recommendations are not hard and fast rules. Most are examples of good programming practices.

For more information about C pointer pragmas, see the *DEC C User's Guide for OpenVMS Systems*.

6.1 Quadword/Longword Argument Pointer Guidelines

Because OpenVMS Alpha 64-bit adressing support allows application programs to access data in 64-bit address spaces, pointers that are not 32-bit sign-extended values (64-bit pointers) will become more common within applications. Existing 32-bit APIs will continue to be supported, and the existence of 64-bit pointers creates some potential pitfalls that programmers must be aware of.

For example, 64-bit addresses may be inadvertently passed to a routine that can handle only a 32-bit address. Another dimension of this would be a new API that includes 64-bit pointers embedded in data structures. Such pointers might be restricted to point to 32-bit address spaces initially, residing within the new data structure as a sign-extended 32-bit value.

Routines should guard against programming errors where 64-bit addresses are being passed instead of 32-bit addresses. This type of checking is called sign-extension checking, which means that the address is checked to ensure that the upper 32 bits are all zeros or all ones, matching the value of bit-31. This checking can be performed at the routine interface that is imposing this restriction.

When defining a new routine interface, you should consider the ease of programming a call to the routine from a 32-bit source module. And you should consider calls written in all OpenVMS programming languages, not just those languages initially supporting 64-bit addressing. To avoid promoting awkward programming practices for the 32-bit caller of a new routine, you should accommodate 32-bit callers as well as 64-bit callers.

Arguments passed by reference that are restricted to reside in a 32-bit address space (P0/P1/S0/S1) should have their reference addresses sign-extension checked.

The OpenVMS Calling Standard requires that 32-bit values passed to a routine be sign-extended to 64-bits before the routine is called. Therefore, the called routine always receives 64-bit values. A 32-bit routine cannot tell if its caller correctly called the routine with a 32-bit address, unless the reference to the argument is checked for sign-extension.

This sign-extension checking would also apply to the reference to a descriptor when data is being passed to a routine by descriptor.

The called routine should return the error status SS$_ARG_GTR_32_BITS if the sign-extension check fails.

Alternately, if you want the called routine to accept the data being passed in a 64-bit location without error and if the sign-extension check fails, the data can be copied by the called routine to a 32-bit address space. The 32-bit address space to which the routine copies the data can be local routine storage (that is, the current stack). If the data is copied to a 32-bit location other than local storage, memory leaks and reentrancy issues must be considered.

When new routines are developed, pointers to code and all data pointers passed to the new routines, should be accommodated in 64-bit address spaces where possible. This is desirable even if the data is a routine or is typically considered "static data", which the programmer, compiler, or linker would not naturally put in a 64-bit address space in this release. When code and static data is supported in 64-bit address spaces, this routine should not need additional changes.

32-bit descriptor arguments should be validated to be 32-bit descriptors.

Routines that accept descriptors should test the fields that allow you to distinguish the 32-bit and 64-bit descriptor forms. If a 64-bit descriptor is received, the routine should return an error.

Most existing 32-bit routines will return or signal the error status SS$_ACCVIO when incorrectly presented with a 64-bit descriptor for the following reasons:

- The 64-bit form of a descriptor contains a MBO (must be one) word at offset 0, where the 32-bit descriptor LENGTH is located, and
- A MBMO (must be minus one) longword at offset 4, where the 32-bit descriptor's POINTER is located as shown in the following figure:

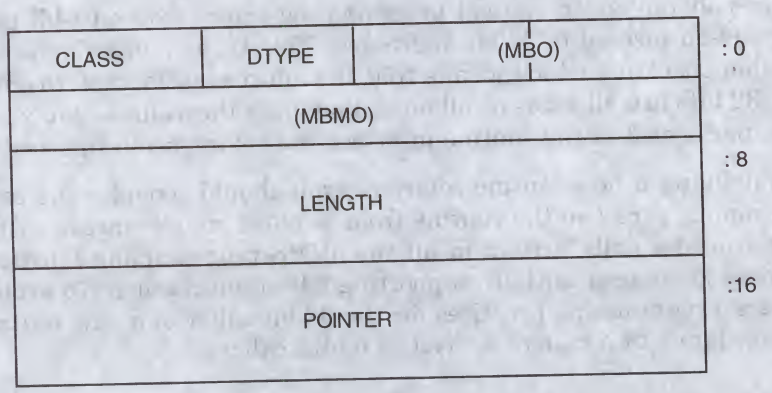

ZK–8489A–GE

Routines that accept arguments passed by 64-bit descriptors should accommodate 32-bit descriptors as well as 64-bit descriptors.

New routines should accommodate 32-bit and 64-bit descriptors within the same routine. The same argument can point to either a 32-bit or 64-bit descriptor.

The 64-bit descriptor MBO word at offset 0 should be tested for one, and the 64-bit descriptor MBMO longword at offset 4 should be tested for a minus one to distinguish between a 64-bit and 32-bit descriptor.

Consider an existing 32-bit routine that is being converted to handle 64-bit as well as 32-bit descriptors. If the input descriptor is determined to be a 64-bit descriptor, the data being pointed to by the 64-bit descriptor can first be copied to a 32-bit memory location, then a 32-bit descriptor would be created in 32-bit memory. This new 32-bit descriptor can then be passed to the existing 32-bit code, so that no further modifications need to be made internally to the routine.

Avoid passing pointers by reference.

If passing a pointer by reference is necessary, as with certain memory management routines, the pointer should be defined to be 64-bit wide.

Mixing 32-bit and 64-bit pointers can cause programming errors when the caller incorrectly passes a 32-bit wide pointer by reference when a 64-bit wide pointer is expected.

If the called routine reads a 64-bit wide pointer that was allocated only a longword by the programmer, the wrong address could be used by the routine.

If the called routine returns a 64-bit pointer, and therefore writes a 64-bit wide address into a longword allocated by the programmer, data corruption can occur.

Existing routines that are passed pointers by reference require new interfaces for 64-bit support. Old routine interfaces would still be passed the pointer in a 32-bit wide memory location and the new routine interface would require that the pointer be passed in a 64-bit wide memory location. Keeping the same interface and passing it 64-bit wide pointers would break existing programs.

Example

The return virtual address used in the new SYS$CRETVA_64 service. Virtual addresses created in P0 and P1 space are guaranteed to have only 32 bits of significance, however all 64 bits are returned. SYS$CRETVA_64 can also create address space in 64-bit space and thus return a 64-bit address. The value that is returned must always be 64 bits because a 64-bit address can be returned.

Memory allocation routines should return the pointer to the data allocated by value (that is, in R0), if possible. The C allocation routines, malloc, calloc, and realloc are examples of this.

New interfaces for routines that are not memory management routines should avoid defining output arguments to receive addresses. Problems will arise whenever a 64-bit subsystem allocates memory and then returns a pointer back to a 32-bit caller in an output argument. The caller may not be able to support or express a 64-bit pointer. Instead of returning a pointer to some data, the caller should provide a pointer to a buffer and the called routine should copy the data into the user's buffer.

A 64-bit pointer passed by reference should be defined in such a way that a call to the routine can be written in a 64-bit language or a 32-bit language. It should be clearly indicated that a 64-bit pointer is required to be passed by all callers.

Routines must not return 64-bit addresses unless they are specifically requested.

It is extremely important that routines which allocate memory and return an address to their callers always allocate 32-bit addressable memory, unless it is known absolutely that the caller is capable of handling 64-bit addresses. This is true for both function return values and output parameters. This rule prevents 64-bit addresses from "creeping in" to applications which do not expect them. As a result, programmers developing callable libraries should be particularly careful to follow this rule.

Suppose an existing routine returns the address of memory it has allocated, as the routine value. If the routine accepts an input parameter which in some way allows it to determine that the caller is 64-bit capable, it is safe to return a 64-bit address. Otherwise, it MUST continue to return 32-bit sign-extended addresses. In the latter case, a new version of the routine could be provided, which 64-bit callers could invoke instead of the existing version if they prefer that 64-bit memory be allocated.

Example: The routines in LIBRTL which manipulate string descriptors can be sure that a caller is 64-bit capable if the descriptor passed in is in the new 64-bit format. As a result, it is safe for them to allocate 64-bit memory for string data, in that case. Otherwise, they will continue to use only 32-bit addressable memory.

Avoid embedded pointers in data structures in public interfaces.

If embedded pointers are necessary for a new structure in a new interface, provide storage within the structure for a 64-bit pointer (quadword aligned). The called routine, which may have to read the pointer from the structure, simply reads all 64 bits.

If the pointer is constrained to be a 32-bit sign-extended address (for example, because the pointer will be passed to a 32-bit routine) a sign-extension check should be performed on the 64-bit pointer at the entrance to the routine. If the sign-extension check fails, the error status SS$_ARG_GTR_32_BITS may be returned to the caller, or the data found to reside in a 64-bit address space may be copied to a 32-bit address space.

The new structure should be defined in such a way that a 64-bit caller or a 32-bit caller does not contain awkward code. The structure should provide a quadword field for the 64-bit caller overlaid with two longword fields for the 32-bit caller. The first of these longwords is the 32-bit pointer field and the next is a MBSE (must be sign-extension) field. For most 32-bit callers, the MBSE field will be zero because the pointer will be a 32-bit process space address. The key here is to define the pointer as a 64-bit value and make it clear to the 32-bit caller that the full quadword must be filled in.

In the following example, both 64-bit and 32-bit callers would pass a pointer to the "block" structure and use the same function prototype when calling the function "routine". (Assume "data" is an unknown structure defined in another module.)

```
#pragma required_pointer_size save
#pragma required_pointer_size 32

typedef struct block {
    int blk_l_size;
    int blk_l_flags;
    union  {
#pragma required_pointer_size 64
        struct data *blk_pq_pointer;
#pragma required_pointer_size 32
        struct  {
            struct data *blk_ps_pointer;
            int blk_l_mbse;
            } blk_r_long_struct;
        } blk_r_pointer_union;
    } BLOCK;

#define blk_pq_pointer     blk_r_pointer_union.blk_pq_pointer
#define blk_r_long_struct  blk_r_pointer_union.blk_r_long_struct
#define blk_ps_pointer     blk_r_long_struct.blk_ps_pointer
#define blk_l_mbse         blk_r_long_struct.blk_l_mbse

/* Routine accepts 64-bit pointer to the "block" structure */
#pragma required_pointer_size 64
int routine(struct block*);

#pragma required_pointer_size restore
```

For an existing 32-bit routine specifying an input argument, which is a structure that embeds a pointer, you can use a different approach to preserve the existing 32-bit interface. You can develop a 64-bit form of the data structure that is distinguished from the 32-bit form of the structure at run-time. Existing code that accepts only the 32-bit form of the structure should automatically fail when presented with the 64-bit form.

The structure definition for the new 64-bit structure should contain the 32-bit form of the structure. Including the 32-bit form of the structure allows the called routine to declare the input argument as a pointer to the 64-bit form of the structure and cleanly handle both cases.

Two different function prototypes can be provided for languages that provide type checking. The default function prototype should specify the argument as a pointer to the 32-bit form of the structure. The 64-bit form of the function prototype can be selected by defining a symbol, specified by documentation.

The 64-bit versus 32-bit descriptor is an example of how this can be done.

Example: In the following example, the state of the symbol FOODEF64 selects the 64-bit form of the structure along with the proper function prototype. If the symbol FOODEF64 is undefined, the old 32-bit structure is defined and the old 32-bit function prototype is used.

The source module that implements the function foo_print would define the symbol FOODEF64 and be able to handle calls from 32-bit and 64-bit callers. The 64-bit caller would set the field foo64$l_mbmo to -1. Foo_print would test the field foo64$l_mbmo for -1 to determine if the caller used the 64-bit form of the structure or the 32-bit form of the structure.

```
#pragma required_pointer_size save
#pragma required_pointer_size 32

typedef struct foo {
    short int     foo$w_flags;
    short int     foo$w_type;
    struct data * foo$ps_pointer;
    } FOO;

#ifndef FOODEF64

/* Routine accepts 32-bit pointer to "foo" structure */
int foo_print(struct foo * foo_ptr);

#endif

#ifdef FOODEF64

typedef struct foo64 {
    union  {
        struct  {
            short int     foo64$w_flags;
            short int     foo64$w_type;
            int           foo64$l_mbmo;
#pragma required_pointer_size 64
            struct data * foo64$pq_pointer;
#pragma required_pointer_size 32
            } foo64$r_foo64_struct;
        FOO foo64$r_foo32;
        } foo64$r_foo_union;
    } FOO64;

#define foo64$w_flags    foo64$r_foo_union.foo64$r_foo64_struct.foo64$w_flags
#define foo64$w_type     foo64$r_foo_union.foo64$r_foo64_struct.foo64$w_type
#define foo64$l_mbmo     foo64$r_foo_union.foo64$r_foo64_struct.foo64$l_mbmo
#define foo64$pq_pointer foo64$r_foo_union.foo64$r_foo64_struct.foo64$pq_pointer
#define foo64$r_foo32    foo64$r_foo_union.foo64$r_foo32

/* Routine accepts 64-bit pointer to "foo64" structure */
#pragma required_pointer_size 64
int foo_print(struct foo64 * foo64_ptr);

#endif

#pragma required_pointer_size restore
```

In the previous example, if the structures "foo" and "foo64" will be used interchangeably within the same source module, you can eliminate the symbol FOODEF64. The routine foo_print would then be defined as follows:

```
int foo_print (void * foo_ptr);
```

Eliminating the FOODEF64 symbol allows 32-bit and 64-bit callers to use the same function prototype, however less strict type checking is then available during the C source compilation.

6.2 Alpha/VAX Guidelines

Only address, size, and length arguments should be passed as quadwords by value.

Arguments passed by value are restricted to longwords on VAX. To be compatible with VAX APIs, quadword arguments should be passed by reference instead of by value. However, addresses, sizes and lengths are examples of arguments which, because of the architecture, could logically be longwords on OpenVMS VAX and quadwords on OpenVMS Alpha.

Even if the API will not be available on OpenVMS VAX, this guideline should still be followed for consistency across all APIs.

Avoid page size dependent units.

Arguments such as lengths and offsets should be represented in units that are page size independent, such as bytes.

A pagelet is an awkward unit. It was invented for compatibility with VAX and is used on OpenVMS Alpha within OpenVMS VAX compatible interfaces. A pagelet is equivalent in size to a VAX page and should not be considered a page size independent unit, because it is often confused with a CPU-specific page on Alpha.

Example: Length_64 argument in EXPREG_64 is passed as a quadword byte count by value.

Naturally align all data passed by reference.

The called routine should specify to the compiler that arguments are aligned, and the compiler can perform more efficient load and store sequences. If the data is not naturally aligned, users will experience performance penalties.

If the called routine can execute incorrectly because the data passed by reference is not naturally aligned, the called routine should do explicit checking and return an error if not aligned. For example, if a load/locked, store/conditional is being done internally in the routine on the data; and the data is not aligned, the load/locked, store/conditional will not work properly.

6.3 Promoting an API from a 32-Bit API to a 64-Bit API

For ease of use, it is best to separate promoting an API from improving the 32-bit design or adding new functionality. Calling a routine within the new 64-bit API should be an easy programming task.

64-bit routines should accept 32-bit forms of structures as well as 64-bit forms.

To make it easy to modify calls to an API, the 32-bit form of a structure should be accepted by the interface as well as the 64-bit form.

Example: If the 32-bit API passed information by descriptor, the new interface should pass the same information by descriptor.

64-bit routines should provide the same functionality as the 32-bit routines.

An application currently calling the 32-bit API should be able to completely upgrade to calling the 64-bit API without having to preserve some of the old calls to the old 32-bit API just because the new 64-bit API is not a functional superset of the old API.

Example: SYS$EXPREG_64 works for P0, P1 and P2 process space. Callers can replace all calls to SYS$EXPREG since SYS$EXPREG_64 is a functional superset of $EXPREG.

Use the suffix "_64" when appropriate.

For system services, this suffix is used for routines that accept 64-bit addresses by reference. For promoted routines, this distinguishes the 64-bit capable version from its 32-bit counterpart, and for new routines, it is a visible reminder that a 64-bit wide address cell will be read/written. This is also used when a structure is passed which contains an embedded 64-bit address, IF the structure is not self-identifying as a 64-bit structure. Hence, a routine name need not include "_64" simply because it receives a 64-bit decriptor. Remember that passing an arbitrary value by reference does not mean the suffix is required; passing a 64-bit address by reference does.

This practice is recommended for other routines as well.

Examples:

SYS$EXPREG_64(region_id_64, length_64, acmode, return_va_64, return_length_64)
SYS$CMKRNL_64(routine_64, quad_arglst_64)

6.4 Example of a 32-Bit Routine and a 64-Bit Routine

The following example illustrates a 32-bit routine interface that has been promoted to support 64-bit addressing. It handles several of the issues addressed in the guidelines.

The C function declaration for an old system service SYS$CRETVA looks like the following:

```
#pragma required_pointer_size save
#pragma required_pointer_size 32
int sys$cretva (
        struct _va_range * inadr,
        struct _va_range * retadr,
        unsigned int       acmode);
#pragma required_pointer_size restore
```

The C function declaration for a new system service SYS$CRETVA_64 looks like the following:

```
#pragma required_pointer_size save
#pragma required_pointer_size 64
int sys$cretva_64 (
        struct _generic_64 * region_id_64,
        void *               start_va_64,
        unsigned __int64     length_64,
        unsigned int         acmode,
        void **              return_va_64,
        unsigned __int64 *   return_length_64);
#pragma required_pointer_size restore
```

The new routine interface for SYS$CRETVA_64 corrects the embedded pointers within the "_va_range" structure, passes the 64-bit region_id_64 argument by reference and passes the 64-bit length_64 argument by value.

7

OpenVMS Alpha Tools and Utilities That Support 64-Bit Addressing

This chapter briefly describes the following OpenVMS Alpha tools that have been enhanced to support 64-bit virtual addressing.

- OpenVMS Debugger
- System-code debugger
- XDELTA
- Watchpoint utility
- SDA
- LIB$ and CVT$ Facilities of the OpenVMS Run-Time Library

7.1 OpenVMS Debugger

On OpenVMS Alpha systems, the Debugger can access the extended memory made available by 64-bit addressing support. You can examine and manipulate data in the complete 64-bit address space.

You can examine a variable as a quadword by using the new option Quad, which is on the Typecast menu of both the Monitor pull-down menu and the Examine dialog box.

The default type for the debugger is longword, which is appropriate for debugging 32-bit applications. It might be advisable to change the default type to quadword for debugging applications that use the 64-bit address space. To do this, the SET TYPE QUADWORD command.

Note that hexadecimal addresses are now 16-digit numbers on Alpha. For example,

```
DBG> EVALUATE/ADDRESS/HEX %hex 000004A0
00000000000004A0
DBG>
```

The debugger supports 32-bit and 64-bit pointers.

For more information about using the OpenVMS Debugger, see the *OpenVMS Debugger Manual*.

7.2 OpenVMS Alpha System-Code Debugger

The OpenVMS Alpha system-code debugger accepts 64-bit addresses and uses full 64-bit addresses to retrieve information.

7.3 Delta/XDelta

XDELTA has always supported 64-bit addressing on OpenVMS Alpha. Quadword display mode displays full quadwords of information. 64-bit address display mode accepts and displays all addresses as 64-bit quantities.

XDELTA has predefined command strings for displaying the contents of the PFN database. With the PFN database layout changes in OpenVMS Alpha Version 7.0, the command strings and the format of the displays has changed accordingly.

For more information about Delta/XDelta, see the *OpenVMS Delta/XDelta Debugger Manual*.

7.4 LIB$ and CVT$ Facilities of the OpenVMS Run-Time Library

For more information about 64-bit addressing support for the LIB$ and CVT$ facilities of the OpenVMS RTL library, refer to the *OpenVMS RTL Library (LIB$) Manual*.

7.5 Watchpoint Utility

The Watchpoint utility is a debugging tool that maintains a history of modifications that are made to a particular location in shared system space by setting watchpoints on 64-bit addresses. It watches any system address, whether in S0, S1, or S2 space.

A $QIO interface to the Watchpoint utility supports 64-bit addresses. The WATCHPOINT command interpreter (WP) issues $QIO request to the WATCHPOINT driver (WPDRIVER) from commands that follow the standard rules of DCL grammar.

Enter commands at the WATCHPOINT> prompt to set, delete, and obtain information from watchpoints. Before invoking the WATCHPOINT command interpreter (WP) or loading the WATCHPOINT driver, you must set the SYSGEN MAXBUF dynamic parameter to 64000, as follows:

```
$ RUN SYS$SYSTEM:SYSGEN
SYSGEN> SET MAXBUF 64000
SYSGEN> WRITE ACTIVE
SYSGEN> EXIT
```

Before invoking WP, you must install the WPDRIVER with SYSMAN, as follows:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> IO CONNECT WPA0/DRIVER=SYS$WPDRIVER/NOADAPTER
SYSMAN> EXIT
```

You can then invoke WP with the following command:

```
$ RUN SYS$SYSTEM:WP
```

Now you can enter commands at the WATCHPOINT> prompt to set, delete, and obtain information from watchpoints.

Yu can best view the WP help screens as well as the output to the Watchpoint utility using a terminal set to 132 characters, as follows:

```
$ SET TERM/WIDTH=132
```

7.6 SDA

As of OpenVMS Alpha Version 7.0, SDA allows a user to specify 64-bit addresses and 64-bit values in expressions. It will also display full 64-bit values where appropriate.

For more information about using SDA 64-bit addressing support, see the *OpenVMS Alpha System Dump Analyzer Utility Manual*.

8 DEC C RTL Support for 64-Bit Addressing

This chapter describes the 64-bit addressing support provided by the DEC C Run-Time Library on OpenVMS Alpha Version 7.0 systems and higher.

The DEC C Run-Time Library includes the following features in support of 64-bit pointers:

- Guaranteed binary and source compatibility of existing programs
- No impact on applications that are not modified to exploit 64-bit support
- Enhanced memory allocation routines that allocate 64-bit memory
- Widened function parameters to accommodate 64-bit pointers
- Dual implementations of functions that need to know the pointer size used by the caller
- New information available to the DEC C Version 5.2 compiler or higher to seamlessly call the correct implementation
- Ability to explicitly call either the 32-bit or 64-bit form of functions for applications that mix pointer sizes
- A single shareable image for use by 32-bit and 64-bit applications

8.1 Using the DEC C Run-Time Library

The DEC C Run-Time Library on OpenVMS Alpha Version 7.0 systems and higher can generate and accept 64-bit pointers. Functions that require a second interface to be used with 64-bit pointers reside in the same object libraries and shareable images as their 32-bit counterparts. No new object libraries or shareable images are introduced. Using 64-bit pointers does not require changes to your link command or link options files.

The DEC C 64-bit environment allows an application to use both 32-bit and 64-bit addresses. For more information about how to manipulate pointer sizes, see the /POINTER_SIZE qualifier and `#pragma pointer_size` and `#pragma required_pointer_size` preprocessor directives in the *DEC C User's Guide for OpenVMS Systems*.

The /POINTER_SIZE qualifier requires you to specify a value of 32 or 64. This value is used as the default pointer size within the compilation unit. As an application programmer, you can compile one set of modules using 32-bit pointers and another set using 64-bit pointers. Use care when these two separate groups of modules call each other.

Use of the /POINTER_SIZE qualifier also influences the processing of DEC C RTL header files. For functions that have a 32-bit and 64-bit implementation, specifying /POINTER_SIZE enables function prototypes to access both functions, regardless of the actual value supplied to the qualifier. In addition, the value

specified to the qualifier determines the default implementation to call during that compilation unit.

The `#pragma pointer_size` and `#pragma required_pointer_size` preprocessor directives can be used to change the pointer size in effect within a compilation unit. You can default pointers to 32-bit pointers and then declare specific pointers within the module as 64-bit pointers. You would also need to specifically call the `_malloc64` form of `malloc` to obtain memory from the 64-bit memory area.

8.2 Obtaining 64-Bit Pointers to Memory

The DEC C RTL has many functions that return pointers to newly allocated memory. In each of these functions, the application owns the memory pointed to and is responsible for freeing that memory.

Functions that allocate memory are:

```
malloc
calloc
realloc
strdup
```

Each of these functions has a 32-bit and 64-bit implementation. When the /POINTER_SIZE qualifier is used, the following functions can also be called:

```
_malloc32, _malloc64
_calloc32, _calloc64
_realloc32, _realloc64
_strdup32, _strdup64
```

When /POINTER_SIZE=32 is specified, all `malloc` calls default to `_malloc32`.

When /POINTER_SIZE=64 is specified, all `malloc` calls default to `_malloc64`.

Regardless of whether the application calls a 32-bit or 64-bit memory allocation routine, there is still a single `free` function. This function accepts either pointer size.

Note that the memory allocation functions are the only ones that return pointers to 64-bit memory. All DEC C RTL structure pointers returned to the calling application (such as a FILE, WINDOW, or DIR) are always 32-bit pointers. This allows both 32-bit and 64-bit callers to pass these structure pointers within the application.

8.3 DEC C Header Files

The header files distributed with DEC C Version 5.2 and higher support 64-bit pointers. Each function prototype whose signature contains a pointer is constructed to indicate the size of the pointer accepted.

A 32-bit pointer can be passed as an argument to functions that accept either a 32-bit or 64-bit pointer for that argument.

A 64-bit pointer, however, cannot be passed as an argument to a function that accepts a 32-bit pointer. Attempts to do this are diagnosed by the compiler with a MAYLOSEDATA message. The diagnostic message IMPLICITFUNC means the compiler can do no additional pointer-size validation for calls to that function.

You might find the following pointer-size compiler diagnostics useful:

- %CC-IMPLICITFUNC

A function prototype was not found before using the specified function. The compiler and run-time system rely on prototype definitions to detect incorrect pointer-size usage. Failure to include the proper header files can lead to incorrect results or pointer truncation.

- %CC-MAYLOSEDATA

  A truncation is necessary to do this operation. The operation could be passing a 64-bit pointer to a function that does not support a 64-bit pointer in the given context. Or it could be a function returning a 64-bit pointer to a calling application that is trying to store that return value in a 32-bit pointer.

- %CC-MAYHIDELOSS

  This message (when enabled) helps expose real MAYLOSEDATA messages that are being suppressed because of a cast operation.

8.4 Functions Affected

The DEC C RTL shipped with OpenVMS Alpha Version 7.0 accommodates applications that use only 32-bit pointers, only 64-bit pointers, or a combination of both. To use 64-bit memory, you must, at a minimum, recompile and relink an application. The amount of source code change required depends on the application itself, calls to other run-time libraries, and the combinations of pointer sizes used.

With respect to 64-bit pointer support, the functions in the DEC C RTL fall into four categories:

- Functions not impacted by choice of pointer size
- Functions enhanced to accept either pointer size
- Functions having a 32-bit and 64-bit implementation
- Functions that accept only 32-bit pointers

From an application developer's perspective, the first two types of functions are the easiest to use in either a single-pointer or mixed-pointer mode.

The third type requires no modifications when used in a single-pointer compilation but might require source code changes when used in a mixed-pointer mode.

The fourth type requires careful attention whenever 64-bit pointers are used.

8.4.1 No Pointer-Size Impact

The choice of pointer-size has no impact on a function if its prototype contains no pointer-related parameters or return values. The mathematical functions are good examples of this.

Even some functions in this category that do have pointers in their prototype are not impacted by pointer size. For example, `strerror` has the prototype:

```
char * strerror (int error_number);
```

This function returns a pointer to a character string, but this string is allocated by the DEC C RTL. As a result, to support both 32-bit and 64-bit applications, these types of pointers are guaranteed to fit in a 32-bit pointer.

8.4.2 Functions Accepting Both Pointer Sizes

The Alpha architecture supports 64-bit pointers. The OpenVMS Alpha Calling Standard specifies that all arguments are actually passed as 64-bit values. Before OpenVMS Alpha Version 7.0, all 32-bit addresses passed to procedures were sign-extended into this 64-bit parameter. The called function declared the parameters as 32-bit addresses, which caused the compiler to generate 32-bit instructions (such as LDL) to manipulate these parameters.

Many functions in the DEC C RTL are enhanced to receive the full 64-bit address. For example, consider `strlen`:

```
size_t strlen (const char *string);
```

The only pointer in this function is the character-string pointer. If the caller passes a 32-bit pointer, the function works with the sign-extended 64-bit address. If the caller passes a 64-bit address, the function works with that address directly.

The DEC C RTL continues to have only a single entry point for functions in this category. There are no source-code changes required to add any of the four pointer-size options for functions of this type. The OpenVMS documentation refers to these functions as 64-bit friendly.

8.4.3 Functions With Two Implementations

There are many reasons why a function might need two implementations—one for 32-bit pointers, the other for 64-bit pointers. Some of these reasons include:

- The pointer size of the return value is the same size as the pointer size of one of the arguments. If the argument is 32 bits, the return value is 32 bits. If the argument is 64bits, the return value is 64 bits.
- One of the arguments is a pointer to an object whose size is pointer-size sensitive. To know how many bytes are being pointed to, the function must know whether the code was compiled in 32-bit or 64-bit pointer-size mode.
- The function returns the address of dynamically allocated memory. The memory is allocated in 32-bit space when compiled for 32-bit pointers and in 64-bit space when compiled for 64-bit pointers.

From the application developer's point of view, there are three function prototypes for each of these functions. The `<string.h>` header file contains many functions whose return value is dependent upon the pointer size used as the first argument to the function call. For example, consider the `memset` function. The header file defines three entry points for this function:

```
void * memset   (void *memory_pointer, int character, size_t size);
void *_memset32 (void *memory_pointer, int character, size_t size);
void *_memset64 (void *memory_pointer, int character, size_t size);
```

The first prototype is the function that your application would currently call if using this function. The compiler changes a call to `memset` into a call to either `_memset32` when compiled /POINTER_SIZE=32, or `_memset64` when compiled /POINTER_SIZE=64.

You can override this default behavior by directly calling either the 32-bit or the 64-bit form of the function. This accommodates applications using mixed pointer sizes, regardless of the default pointer size specified with the /POINTER_SIZE qualifier.

Note that if the application is compiled without specifying the /POINTER_SIZE qualifier, *neither* the 32-bit specific nor the 64-bit specific function prototypes are defined. In this case, the compiler automatically calls the 32-bit interface for all interfaces having dual implementations.

Table 8–1 shows the DEC C RTL functions that have dual implementations in support of 64-bit pointer size. When compiling with the /POINTER_SIZE qualifier, calls to the unmodified function names are changed to calls to the function interface that matches the pointer size specified with the qualifier.

Table 8–1 Functions With Dual Implementations

basename	malloc	strpbrk	wcsncat
bsearch	mbsrtowcs	strptime	wcsncpy
calloc	memccpy	strrchr	wcspbrk
catgets	memchr	strsep	wcsrchr
ctermid	memcpy	strstr	wcsrtombs
cuserid	memmove	strtod	wcsstr
dirname	memset	strtok	wcstok
fgetname	mktemp	strtol	wcstol
fgets	mmap	strtoll	wcstoul
fgetws	qsort	strtoq	wcswcs
fullname	realloc	strtoul	wmemchr
gcvt	rindex	strtoull	wmemcpy
getcap	strcat	strtouq	wmemmove
getcwd	strchr	tgetstr	wmemset
getname	strcpy	tmpnam	
gets	strdup	wcscat	
index	strncat	wcschr	
longname	strncpy	wcscpy	

8.4.4 Functions Restricted to 32-Bit Pointers

Some functions in the DEC C RTL do not support 64-bit pointers. If you try to pass a 64-bit pointer to one of these functions, the compiler generates a %CC-W-MAYLOSEDATA warning. Applications compiled with /POINTER_SIZE=64 might need to be modified to avoid passing 64-bit pointers to these functions.

Table 8–2 shows the functions restricted to using 32-bit pointers. The DEC C RTL offers no 64-bit support for these functions. You must ensure that only 32-bit pointers are used with these functions.

Table 8–2 Functions Restricted to 32-Bit Pointers

atexit	getopt	modf	setstate
execve	iconv	recvmsg	setvbuf
execvp	initstate	sendmsg	
frexp	ioctl	setbuf	

Table 8–3 shows functions that make callbacks to user-supplied functions as part of processing that function call. The callback procedures are not passed 64-bit pointers.

Table 8–3 Callbacks that Pass Only 32-Bit Pointers

from_vms	to_vms
ftw	tputs

8.5 Reading Header Files

This section introduces the pointer-size manipulations used in the DEC C RTL header files. Use the following examples to become more comfortable reading these header files and to help modify your own header files.

Examples

1.
```
:
#if __INITIAL_POINTER_SIZE ❶
#   if (__VMS_VER < 70000000) || !defined __ALPHA ❷
#      error " Pointer size usage not permitted before OpenVMS Alpha V7.0"
#   endif
#   pragma __pointer_size __save ❸
#   pragma __pointer_size 32 ❹
#endif
:
:
#if __INITIAL_POINTER_SIZE ❺
#   pragma __pointer_size 64
#endif
:
:
#if __INITIAL_POINTER_SIZE ❻
#   pragma __pointer_size __restore
#endif
:
```

All DEC C compilers that support the /POINTER_SIZE qualifier predefine the macro __INITIAL_POINTER_SIZE. The DEC C RTL header files take advantage of the ANSI rule that if a macro is not defined, it has an implicit value of 0.

The macro is defined as 32 or 64 when the /POINTER_SIZE qualifier is used. It is defined as 0 if the qualifier is not used. The statement shown as ❶ can be read as "if the user has specified either /POINTER_SIZE=32 or /POINTER_SIZE=64 on the command line".

DEC C Version 5.2 and higher is supported on many OpenVMS platforms. The lines shown as ❷ generate an error message if the target of the compilation is one that does not support 64-bit pointers.

A header file cannot assume anything about the actual pointer-size context in effect at the time the header file is included. Furthermore, the DEC C compiler offers only the __INITIAL_POINTER_SIZE macro and a mechanism to change the pointer size, but no way to determine the current pointer size.

All header files that have a dependency on pointer sizes are responsible for saving ❸, initializing ❹, altering ❺, and restoring ❻ the pointer-size context.

2.
```
:
#ifndef __CHAR_PTR32 ❶
#   define __CHAR_PTR32 1
    typedef char * __char_ptr32;
    typedef const char * __const_char_ptr32;
#endif
:
:
#if __INITIAL_POINTER_SIZE
#   pragma __pointer_size 64
#endif
:
:
#ifndef __CHAR_PTR64 ❷
#   define __CHAR_PTR64 1
    typedef char * __char_ptr64;
    typedef const char * __const_char_ptr64;
#endif
:
```

Some function prototypes need to refer to a 32-bit pointer when in a 64-bit pointer-size context. Other function prototypes need to refer to a 64-bit pointer when in a 32-bit pointer-size context.

DEC C binds the pointer size used in a `typedef` at the time the `typedef` is made. The `typedef` declaration of `__char_ptr32` ❶ is made in a 32-bit context. The `typedef` declaration of `__char_ptr64` ❷ is made in a 64-bit context.

3.
```
:
#if __INITIAL_POINTER_SIZE
#   if (__VMS_VER < 70000000) || !defined __ALPHA
#      error " Pointer size usage not permitted before OpenVMS Alpha V7.0"
#   endif
#   pragma __pointer_size __save
#   pragma __pointer_size 32
#endif
:
❶
:
#if __INITIAL_POINTER_SIZE ❷
#   pragma __pointer_size 64
#endif
:
❸
:
int abs (int __j); ❹
:
__char_ptr32 strerror (int __errnum); ❺
:
```

Before declaring function prototypes that support 64-bit pointers, the pointer context is changed ❷ from 32-bit pointers to 64-bit pointers.

Functions restricted to 32-bit pointers are placed in the 32-bit pointer context section ❶ of the header file. All other functions are placed in the 64-bit context section ❸ of the header file.

Functions that have no pointer-size impact (❹ and ❺) are located in the 64-bit section. Functions that have no pointer-size impact, except for a 32-bit address return value ❺, are also in the 64-bit section, and use the 32-bit specific `typedef`s previously discussed.

4.
```
:
#if __INITIAL_POINTER_SIZE
#   pragma __pointer_size 64
#endif
:
:
#if __INITIAL_POINTER_SIZE == 32 ❶
#   pragma __pointer_size 32
#endif
:
char *strcat (char *__s1, __const_char_ptr64 __s2); ❷
:
#if __INITIAL_POINTER_SIZE
#   pragma __pointer_size 32
    :
    char *_strcat32  (char *__s1, __const_char_ptr64 __s2); ❸
    :
#   pragma __pointer_size 64
    :
    char *_strcat64  (char *__s1, const char *__s2); ❹
    :
#endif
:
```

This example shows declarations of functions that have both a 32-bit and 64-bit implementation. These declarations are located in the 64-bit section of the header file.

The normal interface to the function ❷ is declared using the pointer size specified on the /POINTER_SIZE qualifier. Because the header file is in 64-bit pointer context and because of the statements at ❶, the declaration at ❷ is made using the same pointer size context as the /POINTER_SIZE qualifier.

The 32-bit specific interface ❸ and the 64-bit specific interface ❹ are declared in 32-bit and 64-bit pointer-size context, respectively.

9 MACRO–32 Programming Support for 64-Bit Addressing

This chapter describes the new 64-bit addressing support provided by the MACRO–32 compiler and associated components. The changes are primarily for argument passing and receiving and for address computations.

9.1 Guidelines for 64-Bit Addressing

The following guidelines pertain to using 64-bit addressing in VAX MACRO code that is compiled for OpenVMS Alpha:

- Limit its use to code that you have ported to OpenVMS Alpha.

  For any new development on OpenVMS Alpha, Digital recommends the use of higher-level languages.

- Make 64-bit addressing explicit in your code.

  The 64-bit addressing qualifiers, macros, directives, and built-ins produce code that is more reliable and easier to maintain.

9.2 New and Changed Components for 64-Bit Addressing

The new and changed components that provide MACRO–32 programming support for 64-bit addressing are shown in Table 9–1.

Table 9–1 New and Changed Components for 64-Bit Addressing

Component	Description
$SETUP_CALL64	New macro that initializes the call sequence.
$PUSH_ARG64	New macro that does the equivalent of argument pushes.
$CALL64	New macro that invokes the target routine.
$IS_32BITS	New macro for checking the sign extension of the low 32 bits of a 64-bit value.
$IS_DESC64	New macro for determining if descriptor is a 64-bit format descriptor.
QUAD=NO/YES	New parameter for page macros to support 64-bit virtual addresses.
/ENABLE=QUADWORD	The QUADWORD parameter was extended to include 64-bit address computations.

(continued on next page)

Table 9–1 (Cont.) New and Changed Components for 64-Bit Addressing

Component	Description
.CALL_ENTRY QUAD_ ARGS=TRUE | FALSE	QUAD_ARGS=TRUE | FALSE is a new parameter that indicates the presence (or absence) of quadword references to the argument list.
.ENABLE QUADWORD /.DISABLE QUADWORD	The QUADWORD parameter was extended to include 64-bit address computations.
EVAX_SEXTL	New built-in for sign extending the low 32 bits of a 64-bit value into a destination.
EVAX_CALLG_64	New built-in to support 64-bit calls with variable-size argument lists.
$RAB64 and $RAB64_STORE	New RMS macros for using buffers in 64-bit address space.

9.3 Passing 64-Bit Values

The method that you use for passing 64-bit values depends on whether the size of the argument list is fixed or variable. These methods are described in the following sections.

9.3.1 Calls With a Fixed-Size Argument List

For calls with a fixed-size argument list, use the new macros as shown in Table 9–2.

Table 9–2 Passing 64-Bit Values with a Fixed-Size Argument List

Step	Use...
1. Initialize the call sequence	$SETUP_CALL64
2. "Push" the call arguments	$PUSH_ARG64
3. Invoke the target routine	$CALL64

An example of using these macros follows. Note that the arguments are pushed in reverse order, which is the same way a 32-bit PUSHL instruction is used.

```
MOVL            8(AP), R5       ; fetch a longword to be passed
$SETUP_CALL64   3               ; Specify three arguments in call
$PUSH_ARG64     8(R0)           ; Push argument #3
$PUSH_ARG64     R5              ; Push argument #2
$PUSH_ARG64     #8              ; Push argument #1
$CALL64         some_routine    ; Call the routine
```

The $SETUP_CALL64 macro initializes the state for a 64-bit call. It is required before $PUSH_ARG64 or $CALL64 can be used. If the number of arguments is greater than six, this macro creates a local JSB routine, which is invoked to perform the call. Otherwise, the argument loads and call are inline and very efficient. Note that the argument count specified in the $SETUP_CALL64 does *not* include a pound sign (#). (The standard call sequence requires octaword alignment of the stack with its arguments at the top. The JSB routine facilitates this alignment.)

The inline option can be used to force a call with greater than six arguments to be done without a local JSB routine. However, there are restrictions on its use (see Appendix C).

The $PUSH_ARG64 macro moves the argument directly to the correct argument register or stack location. It is not actually a stack push, but it is the analog of the PUSHL instructions used in a 32-bit call.

The $CALL64 macro sets up the argument count register and invokes the target routine. If a JSB routine was created, it ends the routine. It reports an error if the number of arguments pushed does not match the count specified in $SETUP_CALL64. Both $CALL64 and $PUSH_ARG64 check that $SETUP_CALL64 has been invoked prior to their use.

9.3.1.1 Usage Notes for $SETUP_CALL64, $PUSH_ARG64, and $CALL64

Keep these points in mind when using $SETUP_CALL64, $PUSH_ARG64, and $CALL64:

- The arguments are read as aligned quadwords. To pass a longword from memory, move it to a register first, and then use that register in $PUSH_ARG64, as shown in the example in Section 9.3.1. Similarly, if you know the quadword you want to pass is unaligned, move the value to a register first. Also, keep in mind that indexed operands, such as (R4)[R0], will be evaluated using quadword indexing when used in $PUSH_ARG64.
- If the number of arguments is greater than six, so that a local JSB routine is created, no SP or AP references are allowed between the $SETUP_CALL64 and $CALL64. The $PUSH_ARG64 and $CALL64 macros do report uses of these registers in operands, but they are not allowed in other instructions in this range either and cannot be flagged. To pass an AP- or SP-based argument in this case, move it to a register before the $SETUP_CALL64 invocation.
- If the number of arguments is greater than six, do not rely on values in registers above R15 surviving the $SETUP_CALL64 invocation. Use a nonscratch register as a temporary register instead. For example, suppose you want to pass a value from a stack location, and the call has more than six arguments. In this case, you need to move the value to a register. Rather than using a scratch register such as R28, use a VAX register, such as R0. If all the VAX registers are in use, use R13, R14, or R15.
- It is safe to use the scratch registers above R16 within the range between the $SETUP_CALL64 and the $CALL64. However, you must be careful not to use an argument register that has already been loaded. The argument registers are loaded in downward order, from R21 through R16. So, suppose a call passes six arguments. It is not safe to use R21 after the first $PUSH_ARG64, because that has loaded R21. The $PUSH_ARG64 macro checks for operands that refer to argument registers that have already been loaded. If any are found, the compiler reports a warning. The safest approach is to use registers R22 through R28 when a temporary register is required.

Note

The $SETUP_CALL64, $PUSH_ARG64, and $CALL64 macros are intended to be used in an inline sequence. That is, you cannot branch into the middle of a $SETUP_CALL64/$PUSH_ARG64/$CALL64 sequence, nor can you branch around $PUSH_ARG64 macros or branch out of the sequence to avoid the $CALL64.

For more information about $SETUP_CALL64, $PUSH_ARG64, and $CALL64, see Appendix C.

9.3.2 Calls With a Variable-Size Argument List

For calls with a variable-size argument list, use the new EVAX_CALLG_64 built-in, as shown in the following steps:

1. Create an in-memory argument list.
2. Call a routine, passing the in-memory argument list. For example:

```
EVAX_CALLG_64 (Rn), routine
```

The argument list in the EVAX_CALLG_64 built-in is read as a series of quadwords, beginning with a quadword argument count.

9.4 Declaring 64-Bit Arguments

You can use QUAD_ARGS=TRUE, a new .CALL_ENTRY parameter, to declare the use of quadword arguments in a routine's argument list. With the presence of the QUAD_ARGS parameter, the compiler behaves differently when a quadword reference to the argument list occurs. First, it does not force argument-list homing, which such a reference normally requires. (An argument list containing a quadword value cannot be homed because homing, by definition, packs the arguments into longword slots.) Second, "unaligned memory reference" will not be reported on these quadword references to the argument list.

Note that the actual code generated for the argument-list reference itself is not changed by the presence of the QUAD_ARGS clause, except when the reference is in a VAX quadword instruction, such as MOVQ. For the most part, QUAD_ARGS only prevents argument-list homing due to a quadword reference and suppresses needless alignment messages. This suppression applies to both EVAX_ built-ins and VAX quadword instructions such as MOVQ.

For VAX quadword instructions, the QUAD_ARGS clause causes the compiler to read the quadword argument as it does for EVAX_ built-ins—as an actual quadword. Consider the following example:

```
MOVQ    4(AP), 8(R2)
```

If the QUAD_ARGS clause is specified, MOVQ stores the entire 64 bits of argument 1 into the quadword at 8(R2). If the QUAD_ARGS clause is not specified, MOVQ stores the low longwords of arguments 1 and 2 into the quadword at 8(R2).

Note

Deferred-mode argument list references, such as "@4(AP)", should be avoided in routines that use QUAD_ARGS=TRUE, because the result may not be what you expect. The indirection of this type of reference may require the effective address to be loaded from the argument list if the argument is in memory. This is always performed as a longword load, *not* a quadword load. (If the argument is in a register, it need not be loaded.)

9.4.1 Usage Notes for QUAD_ARGS

Keep these points in mind when using QUAD_ARGS:

- AP-based quadword argument-list references look strange because they appear to overlap. You can improve this situation by defining symbolic names for the argument-list offsets, for example, FIRST_ARG, SECOND_ARG, and so forth. Users are encouraged to define meaningful symbolic names that describe the uses of the arguments to make the source code more readable. Alternatively, you can still use direct argument register references to refer to the first six arguments. Either way, it is useful to declare QUAD_ARGS to ensure that the argument list is not homed.
- Routines that share code must have the same setting for QUAD_ARGS. If they do not, the compiler will report a warning message.
- JSB routines cannot refer to their caller's argument list if the caller has QUAD_ARGS. References to AP within JSB routines require that the last CALL_ENTRY have its argument list homed. HOME_ARGS and QUAD_ARGS are mutually exclusive.
- QUAD_ARGS causes the $ARGn symbols, which the compiler places in the debug symbol table, to be defined as quadwords rather than longwords. These symbols allow easy access to received argument values and can be used in place of register numbers or stack offsets when debugging with the symbolic debugger.

9.5 Specifying 64-Bit Address Arithmetic

There are no explicit pointer-type declarations in MACRO–32. You can create a 64-bit pointer value in a register in a variety of ways. The most common are the EVAX_LDQ built-in for loading an address stored in memory and the MOVA*x* for getting the address of the specified operand.

After a 64-bit pointer value is in a register, an ordinary instruction will access the 64-bit address. The amount of data read from that address depends on the instruction used. Consider the following example:

```
MOVL     4(R1), R0
```

The MOVL instruction reads the longword at offset 4 from R1, regardless of whether R1 contains a 32- or 64-bit pointer.

However, certain addressing modes require the generation of arithmetic instructions to compute the effective address. For VAX compatibility, the compiler computes these as longword operations. For example, `4 + <1@33>` yields the value 4, because the shifted value exceeds 32 bits. If quadword mode is enabled, the upper bit will not be lost.

In compilers shipping with previous versions of OpenVMS Alpha, the /ENABLE=QUADWORD qualifier (and the corresponding .ENABLE QUADWORD and .DISABLE QUADWORD directives) only affected the mode in which constant expression evaluations were performed. For OpenVMS Alpha Version 7.0, these have been extended to affect address computations. They will result in addresses being computed with quadword instructions, such as S*x*ADDQ and ADDQ.

To have quadword operations used throughout a module, specify /ENABLE=QUADWORD on the command line. If you want quadword operations applied only to certain sections, use the .ENABLE QUADWORD and .DISABLE QUADWORD directives to enclose those sections.

There is no performance penalty when using /ENABLE=QUADWORD.

9.5.1 Dependence on Wrapping Behavior of Longword Operations

The compiler cannot use quadword arithmetic for all addressing computations, because existing code may rely on the wrapping behavior of the 32-bit operations. That is, code may perform addressing operations that actually overflow 32 bits, knowing that the upper bits are discarded. Doing the calculation in quadword mode causes an incompatibility.

Before using /ENABLE to set quadword evaluation for an entire module, check the existing code for dependence on longword wrapping. There is no simple way to do this, but as a programming technique, it should be rare and may be called out in the code.

The following example shows the wrapping problem:

```
MOVAL   (R1)[R0], R2
```

Suppose R1 contains the value 7FFFFFFF and R0 contains 1. The MOVAL instruction generates an S4ADDL instruction. The shift and add result exceeds 32 bits, but the stored result is the low 32 bits, sign extended.

If quadword arithmetic were used (S4ADDQ), the true quadword value would result, as shown in the following example:

```
S4ADDL  R0, R1, R2    =>  FFFFFFFF 80000003
S4ADDQ  R0, R1, R2    =>  00000000 80000003
```

The wrapping problem is not limited to indexed-mode addressing. Consider the following example:

```
MOVAB   offset(R1), R0
```

If the symbol offset is not a compile-time constant, this instruction causes a value to be read from the linkage section and added (using an ADDL instruction) to the value in R1. Changing this to ADDQ may change the result if the value exceeds 32 bits.

9.6 Sign Extending and Checking

A new built-in, EVAX_SEXTL (sign-extend longword), is available for sign extending the low 32 bits of a 64-bit value into a destination. This built-in makes explicit the sign extension of the low longword of the source into the destination.

EVAX_SEXTL takes the low 32 bits of the 64-bit value, fills the upper 32 bits with the sign extension (whatever is in bit 31 of the value) and writes the 64-bit result to the destination.

The following examples are all legal uses:

```
evax_sextl r1,r2
evax_sextl r1,(r2)
evax_sextl (r2), (r3)[r4]
```

As shown by these examples, the operands are not required to be registers.

A new macro, $IS_32BITS, is available for checking the sign extension of the low 32 bits of a 64-bit value. It is described in Appendix C.

9.7 Alpha Instruction Built-ins

The compiler supports many Alpha instructions as built-ins. Many of these built-ins (available since the compiler first shipped) can be used to operate on 64-bit quantities. The function of each built-in and its valid operands are documented in *Migrating to an OpenVMS AXP System: Porting VAX MACRO Code*. A full description of each Alpha instruction is documented in the *MACRO–64 Assembler for OpenVMS AXP Systems Reference Manual*.

9.8 Calculating Page-Size Dependent Values

A new parameter, QUAD=NO/YES, for supporting 64-bit virtual addresses is available for each of the page macros, shown in the following list:

- $BYTES_TO_PAGES
- $NEXT_PAGE
- $PAGES_TO_BYTES
- $PREVIOUS_PAGE
- $START_OF_PAGE

These macros provide a standard, architecture-independent means for calculating page-size dependent values. For more information about these macros, see *Migrating to an OpenVMS AXP System: Porting VAX MACRO Code*.

9.9 Creating and Using Buffers in 64-Bit Address Space

The $RAB and $RAB_STORE control block macros have been extended for creating and using data buffers in 64-bit address space. The 64-bit versions are named $RAB64 and $RAB64_STORE, respectively. The rest of the RMS interface is restricted to 32 bits at this time. For more information about $RAB64 and $RAB64_STORE, see Chapter 3.

9.10 Coding for Moves Longer Than 64K Bytes

The MACRO–32 instructions MOVC3 and MOVC5 properly handle 64-bit addresses but the moves are limited to a 64K byte length. This limitation is because MOVC3 and MOVC5 accept word-sized lengths.

For moves longer than 64K bytes, use OTS$MOVE3 and OTS$MOVE5. OTS$MOVE3 and OTS$MOVE5 accept longword-sized lengths. (LIB$MOVC3 and LIB$MOVC5 have the same 64K byte length restriction as MOVC3 and MOVC5.) An example of replacing MOVC3 with OTS$MOVE3 follows.

Code using MOVC3:

```
MOVC3       BUF$W_LENGTH(R5), (R6), OUTPUT(R7)    ; Old code, word length
```

The equivalent 64-bit code with longword length:

```
$SETUP_CALL64  3                    ; Specify three arguments in call
EVAX_ADDQ      R7, #OUTPUT, R7
$PUSH_ARG64    R7                   ; Push destination, arg #3
$PUSH_ARG64    R6                   ; Push source, arg #2
MOVL           BUF$L_LENGTH(R5), R16
$PUSH_ARG64    R16                  ; Push length, arg #1
$CALL64        OTS$MOVE3
```

```
MOVL            BUF$L_LENGTH(R5), R16
EVAX_ADDQ       R6, R16, R1     ; MOVC3 returns address past source
EVAX_ADDQ       R7, R16, R3     ; MOVC3 returns address past destination
```

Because MOVC3 clears R0, R2, R4, and R5, make sure that these side effects are no longer needed.

OTS$MOVE3 and OTS$MOVE5 are documented with other LIBOTS routines in the *OpenVMS RTL General Purpose (OTS$) Manual*.

9.11 Using the MACRO–32 Compiler

In order to take advantage of OpenVMS Alpha 64-bit addressing features, you must use the MACRO–32 compiler included with OpenVMS Alpha Version 7.0.

When you use the latest version of the compiler, regardless of whether you use the 64-bit addressing features, you must also use the latest version of ALPHA$LIBRARY:STARLET.MLB. Make sure that the latest version is installed on your system and that the logical name points to the correct directory.

A

C Macros for 64-Bit Addressing

This appendix describes the following C macros for manipulating 64-bit addresses, for checking the sign extension of the low 32 bits of 64-bit values, and for checking descriptors for the 64-bit format.

- $DESCRIPTOR64
- $is_desc64
- $is_32bits

DESCRIPTOR64

Constructs a 64-bit string descriptor.

Format

$DESCRIPTOR64 name, string

Description

name is Name of variable
string is Address of string

Example:

```
int status;
$DESCRIPTOR64 (gblsec, "GBLSEC_NAME");

...

/* Create global page file section */
status = sys$create_gpfile (&gblsec, 0, 0, section_size, 0, 0);

...
```

This macro resides in descrip.h in SYS$LIBRARY:DECC$RTLDEF.TLB.

$is_desc64

Distinguishes a 64-bit descriptor.

Format

$is_desc64 desc

Description

desc is address of 32-bit or 64-bit descriptor

Returns:

0 if descriptor is 32-bit descriptor
1 if descriptor is 64-bit descriptor

Example:

```
#include <descrip.h>
#include <far_pointers.h>
...
        if ($is_desc64 (user_desc))
        {
                /* Get 64-bit address and 64-bit length from descriptor */
                ...
        }
        else
        {
                /* Get 32-bit address and 16-bit length from descriptor */
                ...
        }
```

This macro resides in descrip.h in SYS$LIBRARY:DECC$RTLDEF.TLB.

$is_32bits

Tests if a quadword is 32-bit sign-extended.

Format

$is_32bits arg

Description

Input: *arg* is 64-bit value

Output:

1 if arg is 32-bit sign-extended
0 if arg is not 32-bit sign-extended

Example:

```
#include <starlet_bigpage.h>
...
if ($is_32bits(user_va))
        counter_32++;   /* Count number of 32-bit references */
else
        counter_64++;  /* Count number of 64-bit references */
```

This macro resides in starlet_bigpage.h in SYS$LIBRARY:SYS$STARLET_C.TLB.

B

64-Bit Example Program

This example program demonstrates the 64-bit region creation and deletion system services. It uses SYS$CREATE_REGION_64 to create a region and then uses SYS$EXPREG_64 to allocate virtual addresses within that region. The virtual address space and the region are deleted by calling SYS$DELETE_REGION_64.

```
/*
 ******************************************************************************
 *
 * Copyright © Digital Equipment Corporation, 1995 All Rights Reserved.
 * Unpublished rights reserved under the copyright laws of the United States.
 *
 * The software contained on this media is proprietary to and embodies the
 * confidential technology of Digital Equipment Corporation.  Possession, use,
 * duplication or dissemination of the software and media is authorized only
 * pursuant to a valid written license from Digital Equipment Corporation.
 *
 * RESTRICTED RIGHTS LEGEND   Use, duplication, or disclosure by the U.S.
 * Government is subject to restrictions as set forth in Subparagraph
 * (c)(1)(ii) of DFARS 252.227-7013, or in FAR 52.227-19, as applicable.
 *
 ******************************************************************************
/*
    This program creates a region in P2 space using the region creation
    service and then creates VAs within that region.  The intent is to
    demonstrate the use of the region services and how to allocate virtual
    addresses within a region.  The program also makes use of 64-bit
    descriptors and uses them to format return values into messages with the
    aid of SYS$GETMSG.

    To build and run this program type:

    $ CC/POINTER_SIZE=32/STANDARD=RELAXED/DEFINE=(__NEW_STARLET=1) -
        REGIONS.C
    $ LINK REGIONS.OBJ
    $ RUN REGIONS.EXE
*/

#include    <descrip.h>          /* Descriptor Definitions                */
#include    <far_pointers.h>     /* Long Pointer Definitions              */
#include    <gen64def.h>         /* Generic 64-bit Data Type Definition   */
#include    <iledef.h>           /* Item List Entry Definitions           */
#include    <ints.h>             /* Various Integer Typedefs              */
#include    <iosbdef.h>          /* I/O Status Block Definition           */
#include    <psldef.h>           /* PSL$ Constants                        */
#include    <ssdef.h>            /* SS$_ Message Codes                    */
#include    <starlet.h>          /* System Service Prototypes             */
#include    <stdio.h>            /* printf                                */
#include    <stdlib.h>           /* malloc, free                          */
#include    <string.h>           /* memset                                */
#include    <syidef.h>           /* $GETSYI Item Code Definitions         */
#include    <vadef.h>            /* VA Creation Flags and Constants       */
```

64-Bit Example Program

```
                                                                    */
/*  Module-wide constants and macros.

#define     BUFFER_SIZE         132
#define     HW_NAME_LENGTH       32
#define     PAGELET_SIZE        512
#define     REGION_SIZE         128

#define     good_status(code)   ((code) & 1)

                                                                    */
/*  Module-wide Variables

int
    page_size;

$DESCRIPTOR64 (msgdsc, "");

                                                                    */
/*  Function Prototypes

int get_page_size (void);
static void print_message (int code, char *string);

main (int argc, char **argv)
{
    int
        i,
        status;

    uint64
        length_64,
        master_length_64,
        return_length_64;

    GENERIC_64
        region_id_64;

    VOID_PQ
        master_va_64,
        return_va_64;

                                                                    */
 /*  Get system page size, using SYS$GETSYI.

    status = get_page_size ();
    if (!good_status (status))
        return (status);

                                                                    */
 /*  Get a buffer for the message descriptor.

    msgdsc.dsc64$pq_pointer = malloc (BUFFER_SIZE);
    printf ("Message Buffer Address = %016LX\n\n", msgdsc.dsc64$pq_pointer);

                                                                    */
 /*  Create a region in P2 space.

    length_64 = REGION_SIZE*page_size;
    status = sys$create_region_64 (
        length_64,                  /* Size of Region to Create        */
        VA$C_REGION_UCREATE_UOWN,   /* Protection on Region            */
        0,                          /* Allocate in Region to Higher VAs */
        &region_id_64,              /* Region ID                       */
        &master_va_64,              /* Starting VA in Region Created   */
        &master_length_64);         /* Size of Region Created          */
    if (!good_status (status))
    {
        print_message (status, "SYS$CREATE_REGION_64");
        return (status);
    }

    printf ("\nSYS$CREATE_REGION_64 Created this Region:  %016LX - %016LX\n",
        master_va_64,
        (uint64) master_va_64 + master_length_64 - 1);
```

```
    /*  Create virtual address space within the region.                    */
        for (i = 0; i < 3; ++i)
        {
            status = sys$expreg_64 (
                &region_id_64,      /* Region to Create VAs In              */
                page_size,          /* Number of Bytes to Create            */
                PSL$C_USER,         /* Access Mode                          */
                0,                  /* Creation Flags                       */
                &return_va_64,      /* Starting VA in Range Created         */
                &return_length_64); /* Number of Bytes Created              */
            if (!good_status (status))
            {
                print_message (status, "SYS$EXPREG_64");
                return status;
            }
            printf ("Filling %016LX - %16LX with %0ds.\n",
                return_va_64,
                (uint64) return_va_64 + return_length_64 - 1,
                i);
            memset (return_va_64, i, page_size);
        }

    /*  Return the virtual addresses created within the region, as well as the
        region itself.                                                      */
        printf ("\nReturning Master Region:  %016LX - %016LX\n",
            master_va_64,
            (uint64) master_va_64 + master_length_64 - 1);

        status = sys$delete_region_64 (
            &region_id_64,      /* Region to Delete                        */
            PSL$C_USER,         /* Access Mode                             */
            &return_va_64,      /* VA Deleted                              */
            &return_length_64); /* Length Deleted                          */

        if (good_status (status))
            printf ("SYS$DELETE_REGION_64 Deleted VAs Between:  %016LX - %016LX\n",
                return_va_64,
                (uint64) return_va_64 + return_length_64 - 1);
        else
        {
            print_message (status, "SYS$DELTE_REGION_64");
            return (status);
        }

    /*  Return message buffer.                                              */
        free (msgdsc.dsc64$pq_pointer);
    }

    /*  This routine obtains the system page size using SYS$GETSYI.  The return
        value is recorded in the module-wide location, page_size.           */
    int get_page_size ()
    {
    int
        status;

    IOSB
        iosb;

    ILE3
        item_list [2];

    /* Fill in SYI item list to retrieve the system page size.              */
```

```
    item_list[0].ile3$w_length       = sizeof (int);
    item_list[0].ile3$w_code         = SYI$_PAGE_SIZE;
    item_list[0].ile3$ps_bufaddr     = &page_size;
    item_list[0].ile3$ps_retlen_addr = 0;
    item_list[1].ile3$w_length       = 0;
    item_list[1].ile3$w_code         = 0;

/* Get the system page size.                                              */

    status = sys$getsyiw (                          /* EFN                */
                0,                                  /* CSI address        */
                0,                                  /* Node name          */
                0,                                  /* Item list          */
                &item_list,                         /* I/O status block   */
                &iosb,                              /* AST address        */
                0,                                  /* AST parameter      */
                0);

    if (!good_status (status))
    {
        print_message (status, "SYS$GETJPIW");
        return (status);
    }
    if (!good_status (iosb.iosb$w_status))
    {
        print_message (iosb.iosb$w_status, "SYS$GETJPIW IOSB");
        return (iosb.iosb$w_status);
    }

    return SS$_NORMAL;
}

/*  This routine takes the message code passed to the routine and then uses
    SYS$GETMSG to obtain the associated message text.  That message is then
    printed to stdio along with a user-supplied text string.               */

#pragma inline (print_message)
static void print_message (int code, char *string)
{
    msgdsc.dsc64$q_length = BUFFER_SIZE;
    sys$getmsg (
        code,                                       /* Message Code       */
        (unsigned short *) &msgdsc.dsc64$q_length,  /* Returned Length    */
        &msgdsc,                                    /* Message Descriptor */
        15,                                         /* Message Flags      */
        0);                                         /* Optional Parameter */
    *(msgdsc.dsc64$pq_pointer+msgdsc.dsc64$q_length) = '\0';
     printf ("Call to %s returned:  %s\n",
        string,
        msgdsc.dsc64$pq_pointer);
}
```

C

MACRO-32 Macros for 64-Bit Addressing

This appendix describes the MACRO-32 macros for manipulating 64-bit addresses, for checking the sign extension of the low 32 bits of 64-bit values, and for checking descriptors for the 64-bit format.

These macros reside in the directory ALPHA$LIBRARY:STARLET.MLB (generally synonymous with SYS$LIBRARY:STARLET.MLB) and can be used by both application code and system code. The page macros have also been enhanced for 64-bit addresses. The support is provided by a new parameter, QUAD=NO/YES.

Note that you can use certain arguments to the macros described in this appendix to indicate register sets. To express a register set, list the registers, separated by commas, within angle brackets. For example:

```
<R1,R2,R3>
```

If the set contains only one register, the angle brackets are not required.

C.1 Macros for Manipulating 64-Bit Addresses

This section describes the following macros, designed to manipulate 64-bit addresses:

- $SETUP_CALL64
- $PUSH_ARG64
- $CALL64

$SETUP_CALL64

Initializes the call sequence.

Format

$SETUP_CALL64 arg_count, inline=true or false

Parameters

arg_count
The number of arguments in the call.

inline
Forces inline expansion, rather than creation of a JSB routine, when set to TRUE. If there are six or fewer arguments, the default is INLINE=FALSE.

Description

This macro initializes the state for a 64-bit call. It *must* be used before using $PUSH_ARG64 and $CALL64.

If there are six or fewer arguments, the code is always in line.

By default, if there are more than six arguments, this macro creates a JSB routine that is invoked to perform the actual call. However, if the inline option is specified as INLINE=TRUE, the code is generated in line. This option should be enabled *only* if the code in which it appears has a fixed stack depth. A fixed stack depth can be assumed if no RUNTIMSTK or VARSIZSTK messages have been reported. Otherwise, if the stack alignment is not at least quadword, there might be many alignment faults in the called routine and in anything the called routine calls. The default behavior (INLINE=FALSE) does not have this problem.

If there are more than six arguments, there can be no references to AP or SP between a $SETUP_CALL64 and the matching $CALL64, because the $CALL64 code may be in a separate JSB routine. In addition, temporary registers (R16 and above) may not survive the $SETUP_CALL64. However, they can be used *within* the range, except where R16 through R21 interfere with the argument registers already set up. In such cases, higher temporary registers should be used instead.

Note

The $SETUP_CALL64, $PUSH_ARG64, and $CALL64 macros are intended to be used in an inline sequence. That is, you cannot branch into the middle of a $SETUP_CALL64/$PUSH_ARG64/$CALL64 sequence, nor can you branch around $PUSH_ARG64 macros or branch out of the sequence to avoid the $CALL64.

$PUSH_ARG64

Does the equivalent of argument pushes for a call.

Format

$PUSH_ARG64 argument

Parameters

argument
The argument to be pushed.

Description

This macro pushes a 64-bit argument for a 64-bit call. The macro $SETUP_CALL64 *must* be used before you can use $PUSH_ARG64.

Arguments will be read as aligned quadwords. That is, $PUSH_ARG64 4(R0) will read the quadword at 4(R0), and push the quadword. Any indexed operations will be done in quadword mode.

To push a longword value from memory as a quadword, first move it into a register with a longword instruction, and then use $PUSH_ARG64 on the register. Similarly, to push a quadword value that you know is *not* aligned, move it to a temporary register first, and then use $PUSH_ARG64.

If the call contains more than six arguments, this macro checks for SP or AP references in the argument. If the call contains more than six arguments, SP references are not allowed, and AP references are allowed only if the inline option is used.

The macro also checks for references to argument registers that have already been set up for the current $CALL64. If it finds such references, a warning is reported to advise the user to be careful not to overwrite an argument register before it is used as the source in a $PUSH_ARG64.

The same checking is done for AP references when there are six or fewer arguments; they are allowed, but the compiler cannot prevent you from overwriting one before you use it. Therefore, if such references are found, an informational message is reported.

Note that if the operand uses a symbol whose name includes one of the strings R16 through R21, *not* as a register reference, this macro might report a spurious error. For example, if the invocation `$PUSH_ARG64 SAVED_R21` is made after R21 has been set up, this macro will unnecessarily report an informational message about overwriting argument registers.

Also note that $PUSH_ARG64 *cannot* be in conditional code. $PUSH_ARG64 updates several symbols, such as the remaining argument count. Attempting to write code that branches around a $PUSH_ARG64 in the middle of a $SETUP_CALL64/$CALL64 sequence will not work properly.

$CALL64

Invokes the target routine.

Format

$CALL64 call_target

Parameters

call_target
The routine to be invoked.

Description

This macro calls the specified routine, assuming $SETUP_CALL64 has been used to specify the argument count, and $PUSH_ARG64 has been used to push the quadword arguments. This macro checks that the number of pushes matches what was specified in the setup call.

The call_target operand must not be AP- or SP-based.

C.2 Macros for Checking Sign Extension and Descriptor Format

The macros in this section are used for checking certain values and directing program flow based on the outcome of the check.

$IS_32BITS

Checks the sign extension of the low 32 bits of a 64-bit value and directs the program flow based on the outcome of the check.

Format

```
$IS_32BITS   quad_arg, leq_32bits, gtr_32bits, temp_reg=22
```

Parameters

quad_arg
A 64-bit quantity, either in a register or in an aligned quadword memory location.

leq_32bits
Label to branch to if quad_arg is a 32-bit sign-extended value.

gtr_32bits
Label to branch to if quad_arg is greater than 32 bits.

temp_reg=22
Register to use as a temporary register for holding the low longword of the source value—R22 is the default.

Description

$IS_32BITS checks the sign extension of the low 32 bits of a 64-bit value and directs the program flow based on the outcome of the check.

Examples

1.
```
$is_32bits  R9, 10$
```

   In this example, the compiler checks the sign extension of the low 32 bits of the 64-bit value at R9 using the default temporary register, R22. Depending on the type of branch and the outcome of the test, the program either branches or continues in line.

2.
```
$is_32bits  4(R8), 20$, 30$, R28
```

   In this example, the compiler checks the sign extension of the low 32 bits of the 64-bit value at 4(R8) using R28 as a temporary register and, based on the check, branches to either 20$ or 30$.

$IS_DESC64

Tests the specified descriptor to determine if it is a 64-bit format descriptor, and directs the program flow based on the outcome of the test.

Format

$IS_DESC desc_addr, target, size=long or quad

Parameters

desc_addr
The address of the descriptor to test.

target
The label to branch to if the descriptor is in 64-bit format.

size=long
The size of the address pointing to the descriptor. Acceptable values are "long" (the default) and "quad".

Description

$IS_DESC64 tests the fields which distinguish a 64-bit descriptor from a 32-bit descriptor. If it is in 64-bit form, a branch is taken to the specified target. The address to be tested is read as a longword, unless SIZE=QUAD is specified.

Examples

1. `$is_desc64 r9, 10$`

   In this example, the descriptor pointed to by R9 is tested, and if it is in 64-bit form, a branch to 10$ is taken.

2. `$is_desc64 8(r0), 20$, size=quad`

   In this example, the quadword at 8(R0) is read, and the descriptor it points to is tested. If it is in 64-bit form, a branch to 20$ is taken.

Index

Q

R

S

V

AA-QSBCA-TE

Printed in Europe

Printed on
Recycled paper